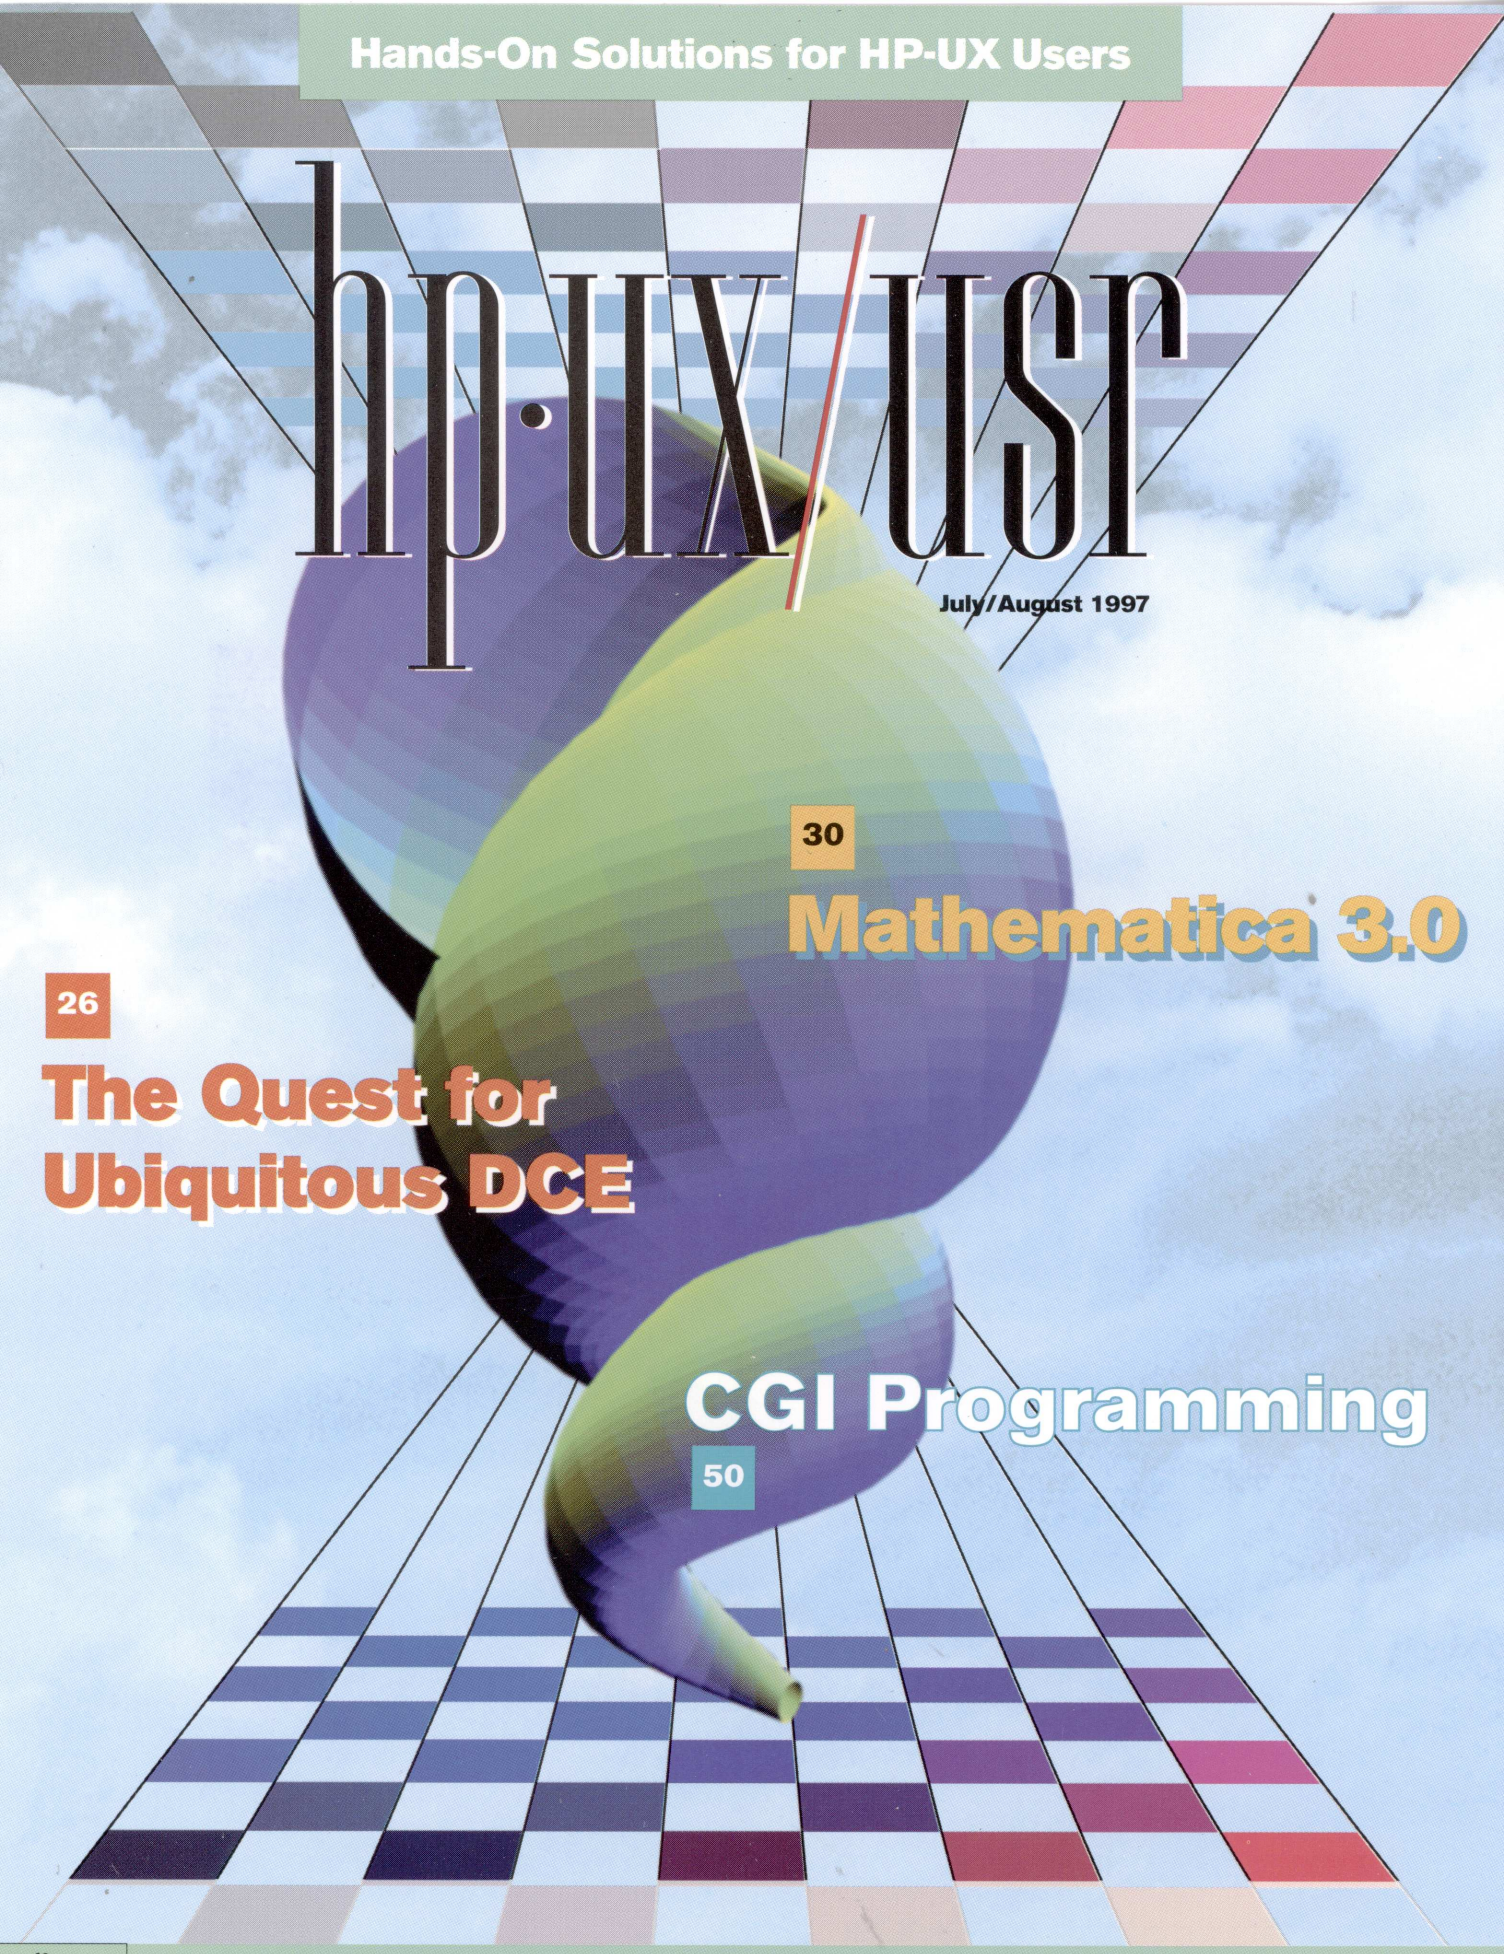
Hands-On Solutions for HP-UX Users

hp·ux/usr

July/August 1997



30 Mathematica 3.0

26 The Quest for Ubiquitous DCE

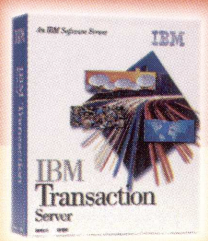
CGI Programming 50

BULK RATE
US Postage
PAID
Senatobia, MS
Permit #30

A Publication of Interex • The International Association of Hewlett-Packard Computing Professionals

IBM





IBM software for HP-UX

Perform extraordinary feats.

Young though she is, "Bending Betty" Bixby tries to look at life from every possible angle. Yet here's an even more unexpected twist. IBM's Transaction Server for HP-UX is so flexible it provides industrial-strength, online transaction processing and transaction management for mission-critical applications not only on your local area network, but also on your wide area network. And its scalability is plenty pliant too – from tens to thousands of users. Transaction server is just one of many powerful and proven IBM solutions for HP-UX. Others include DB2, Net.Data, MQSeries, ADSM, Lotus Domino and more. To get further information – including opportunities for resellers – contact your HP distributor or visit the IBM Software for HP-UX Website at **www.software.ibm.com/hp**

CIRCLE 159 ON READER SERVICE CARD

IBM, MQSeries, DB2, Net.Data, ADSM (ADSTAR Distributed Storage Manager), and Transaction Server are registered trademarks of International Business Machines Corporation. Lotus Domino is a registered trademark of Lotus Corporation. All other company and/or product names are trademarks of their respective companies. ©1997 IBM Corp. All rights reserved.

Is It New Or Is It Refurbished



The only difference is the money you save.



With 17 years of knowledgeable experience and reliable service, Monterey Bay Communications is a leader in Hewlett-Packard workstation remarketing. We're professionals at providing HP 1000 and 9000 users with reliable equipment that is functionally and cosmetically identical to what HP offers — and at substantial cost savings. In addition to the 700 / 400 / 300 / 200 series, Monterey Bay Communications also offers mass storage systems, monitors, memory and interfaces, as well as a variety of printers.

All equipment is warranteed and eligible for Hewlett-Packard maintenance. An extensive parts and spares inventory and knowledgeable staff ensure prompt service and immediate delivery.

For more detailed information or a price quotation, give MBC a call at 408/429-6144.

	NEW	MBC
Performance	✓	✓
Warranty	✓	✓
Support / Tech Expertise	✓	✓
Accessories	✓	✓
Maintenance Eligibility	✓	✓
Substantial Cost Savings		✓
Simple Order Processing		✓
Immediate Delivery		✓

CIRCLE 7 ON READER SERVICE CARD

MONTEREY BAY COMMUNICATIONS

The HP Workstation Remarketing Specialists

Monterey Bay Communications Inc., 1010 Fair Avenue, Santa Cruz, CA 95060 Tel: 408-429-6144 Fax: 408-429-1918

We'd like to thank Hewlett-Packard.

See us in booth 1725 at HP World



"Crystal" Award for significant achievement in HP OpenView Systems and Network Management.

We deeply appreciate the honor and we value our excellent relationship with Hewlett-Packard.

Come to our booth at HP World to see the award, and discover why we earned it!

inoTech

Systems & Network Management

1-800-INOTECH
466-8324

HEWLETT PACKARD
Channel Partner

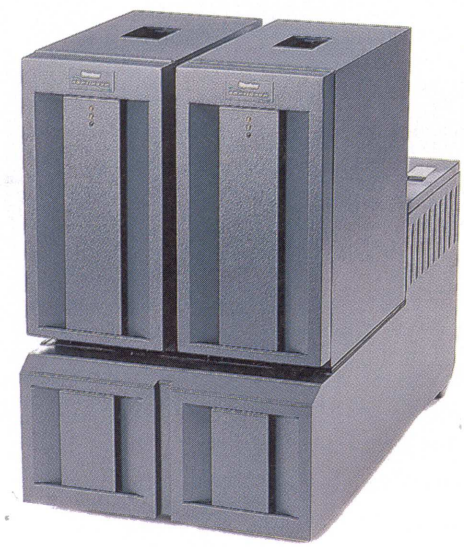
www.inotech.com

CIRCLE 54 ON READER SERVICE CARD



HP-UX™ APPLICATIONS & FAULT TOLERANCE?

CONTINUUM ENTERPRISE SERVER



Stratus and Continuum are registered trademarks of Stratus Computer, Inc. HP-UX is a trademark of the Hewlett-Packard Company.

HOW MUCH WORK WILL THAT BE?

NONE!

No additional hardware. No software. No training. No consulting fees. Now your critical applications can be fault tolerant simply by running them on our platform! The Continuum® Enterprise Server is completely binary-compatible with your HP-UX applications. No recompiling or extra code needed. **Guaranteed!**

Stratus®
The Availability Company

See this working in Booth #403 at HP World '97 – Chicago.

http://www.stratus.com
1(800) 486-9194

CIRCLE 141 ON READER SERVICE CARD

Come to HP WORLD '97 BOOTH 1632 And See How

Secure Single Sign-On With SnareWorks™

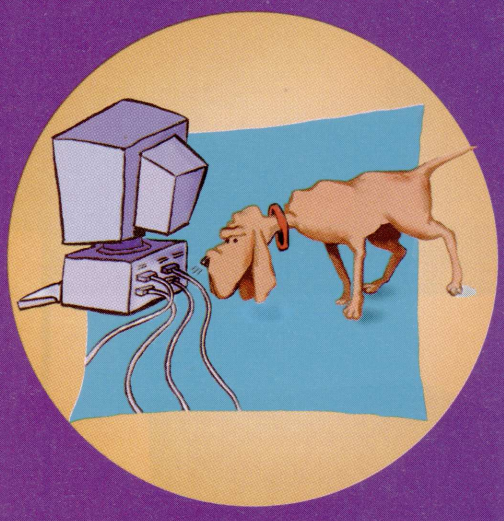
ONE account, **ONE** password…
Is it **REALLY** possible?

With SnareWorks from IntelliSoft it is. By utilizing both secret key and public key encryption technology, combined with SnareWork's transparent operation, you can implement **single sign-on across the enterprise.** All of your intranet applications, the off-the-shelf ones and the ones that you developed in-house, can now be accessible through a single, secure identification and authentication mechanism. And to make matters even more secure, you can configure **any** user to identify themselves through a **password,** a **hardware token, a public key certificate,** or even a **smartcard.**

Think about it.

Or maybe just try it for yourself.

Visit our Web site at **www.isoft.com** to learn more about SnareWorks and **Download a FREE Working Demo!**



FEATURES

- Single Sign-On
- Scalable Authorization Framework
- Centralized Graphical Rule-based Configuration
- Transparent Operation
- No Programming Required
- No Application/Library Replacements
- Extensible to any Application Using Protocol Support Modules™ (PSM)
- Transparent Tunneling for Virtual Private Networks
- Dynamic User Registration for Extranets
- Downloadable Client
- ZERO Management on the Desktop

IntelliSoft
Distributed Computing Solutions

For Information Call 508-635-9070

IntelliSoft Corporation 5 Broadview Road, Acton, MA 01720
Fax: 508-635-9210 URL: http://www.isoft.com E-mail: info@isoft.com

©1997 IntelliSoft Corporation. All rights reserved. SnareWorks and the IntelliSoft logo are trademarks of IntelliSoft Corp. Snare may be a trademark or registered trademark of its respective company.

Portions of SnareWorks are licensed from Snare Networks Corp. snare™ you're secure

CIRCLE 55 ON READER SERVICE CARD

hp·ux/usr

Hands-On Solutions for HP-UX Users

Contents

Volume 5, Number 4 July/August 1997

Columns

22 **HP-UX Systems Administration**
Chris Curtin

66 **HP-UX**
David L. Totsch

68 **Windows NT**
Bob Combs

74 **Internet Goodies**
Joseph Berry

78 **X-Watch**
Larry Headlund

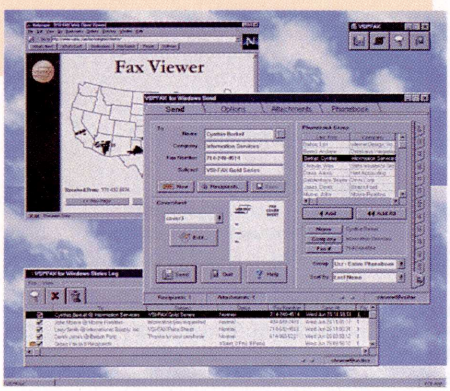
82 **HP-RT**
Anita Harris

Departments

8 **New Products**

16 **Q & A**

88 **CSL Perspective**



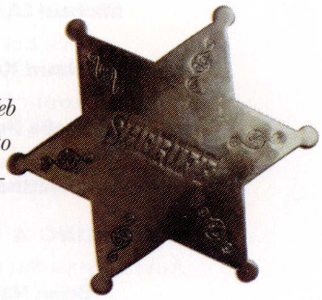
New Products 8

Features

26
The Quest for Ubiquitous DCE

DCE and the Web—a match made in heaven? The DCE/Web Advanced Technology Offering from the Open Group promises to bring to the Web a scaleable security model with built-in authorization. And the Web highlights a broader market that DCE can tackle: the corporate intranet. This article offers an overview of DCE and a glimpse of IntelliSoft's DCE/Snare.

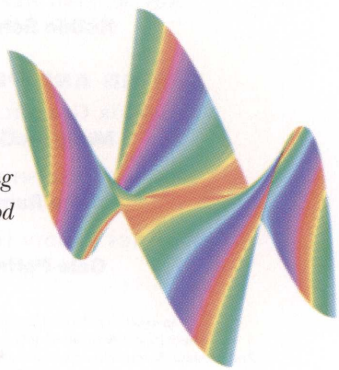
Jonathan Chinitz



30
Mathematica 3.0

Using your workstation for engineering, science, statistics, data analysis, mathematics, or technical education? Here is an introductory look at the latest version of Mathematica from Wolfram Research. If you have been thinking about investigating this vast and flexible program, this article should give you a good idea of what you can do with Mathematica 3.0.

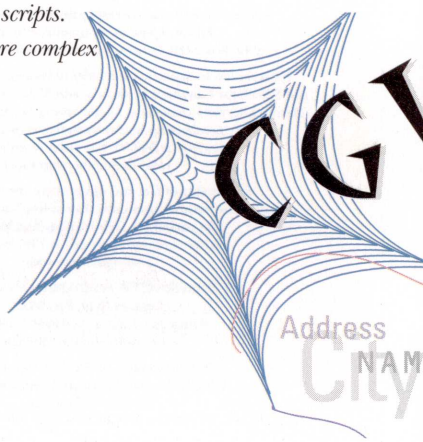
Michael Ehrhardt



50
CGI Programming: Getting the Most from Web-Based Forms

JavaBeans and ActiveX get plenty of hype these days, but you can get a lot out of "simple" CGI scripts. This article—for the Perl and CGI savvy—fills in a glaring gap in CGI coverage available in programming books: treating Web forms as single CGI scripts. Once you grasp this idea, you'll be able to create much more complex forms without the accompanying design problems.

Chris Cobb



Cover Illustration by Michael Ehrhardt

hp-ux/usr
A Publication of **Interex®**

http://www.interex.org/

INTEREX EXECUTIVE DIRECTOR
Charles A. Piercey

PUBLISHER
Michael Doyle

PUBLICATIONS MANAGER
Connie Wright

EDITORIAL

MANAGING EDITOR
Michael Ehrhardt

COPY EDITOR
Richard Kranz

NEW PRODUCTS EDITOR
Michelle Pollace

PROOFREADER
Jean Nattkemper

ADVERTISING & MARKETING

ADVERTISING SALES MANAGER
Brian Hallin

ADVERTISING REPRESENTATIVE
Phil Nguyen

ADVERTISING REPRESENTATIVE
Kathie Schwartz

DESIGN AND PRODUCTION

SENIOR GRAPHICS DESIGNER
Molly McGinnity

DESKTOP PUBLISHING SPECIALIST
June Ramirez

GRAPHICS PRODUCTION SPECIALIST
Gale Patterson

hp-ux/usr is published bimonthly by Interex, the International Association of Hewlett-Packard Computing Professionals. Second-class application pending at Sunnyvale, California 94086 and additional offices. The editorial and business offices are located at 1192 Borregas Ave., Sunnyvale, California 94089, USA, 408.747.0227, Fax 408.747.0947. Address membership questions and change of address to Membership Services. Address all questions concerning circulation/distribution to the Distribution Manager.

Remittances should be sent to Interex, File No. 61054, P. O. Box 60000, San Francisco, California 94160, USA.

Address all editorial correspondence to Michael Ehrhardt, Editor, *hp-ux/usr* Magazine, c/o Interex, P.O. Box 3439, Sunnyvale, California 94088-3439, USA.

Subscription to *hp-ux/usr* is $49.50 (6 issues) per year in the U.S., add $25.00 for Canada and Mexico, add $50.00 for all other countries. Member Services (Associate, Contributing or Online Service Package membership levels) include a subscription to *hp-ux/usr*. at $49.50. For other Member Services refer to membership form.

Statements of fact and opinion are the responsibility of the authors alone and do not imply an opinion on the part of the Interex Board or Magazine. Entire contents copyright © 1997 by Interex. All rights reserved.

POSTMASTER: Send address changes to: Interex, P.O. Box 3439, Sunnyvale, California 94088-3439 USA. Attention: Member Services.

TRADEMARKS: UNIX, X/Open Co. Ltd.; HP-UX, Hewlett-Packard; X Window System, X Consortium, Inc.

How to Contact Interex...

EDITORIAL SUBMISSIONS

hp-ux/usr encourages readers to contribute their opinions, tips, and solutions. When sending letters for publication or to request author contribution guidelines, please address them to *hp-ux/usr* editor Michael Ehrhardt.

Postal Address:
Interex
P.O. Box 3439
Sunnyvale, CA 94088-3439

Office Address:
1192 Borregas Avenue
Sunnyvale, CA 94089

Because of the difference in zip codes between our office address and P.O. Box, please be sure to address all regular mail to the P.O. Box. Any express service packages should be delivered to the Borregas Avenue address. Thank you for your attention to this small but significant detail.

TELEPHONE:

The Interex switchboard is open 8:00 a.m.–5:00 p.m., Pacific Time. Call 800.468.3739 (U.S. and Canada) or 408.747.0227. After 5:30 p.m. our voicemail system will record your call.

FACSIMILE:

Call 408.747.0947

INTERNET:

To send e-mail to Interex, use the following address format:
<IDname>@interex.org
The following IDs are currently active on Interex's HP 9000:

Address	Department
ehrhardt	*hp-ux/usr* Letters to the Editor, Q&A, and requests for author guidelines
webmaster	Internet support
cslhpux	Contributed Software Library
membership	Membership/subscription inquiries and services
pubs	Circulation and advertising inquiries
conference	Conference questions and arrangements

An example would be pubs@interex.org
Anything before the @ sign is case insensitive.

COMPUSERVE:

Interex can be contacted via the CompuServe ↔ Internet gateway. To send CompuServe mail, use the following format:

>INTERNET: pubs@interex.org

You can address your mail to specific departments using the ID's listed above.
Interex maintains a CompuServe account that is collected daily. Please address all messages to ID no. 76376, 1222.

interex

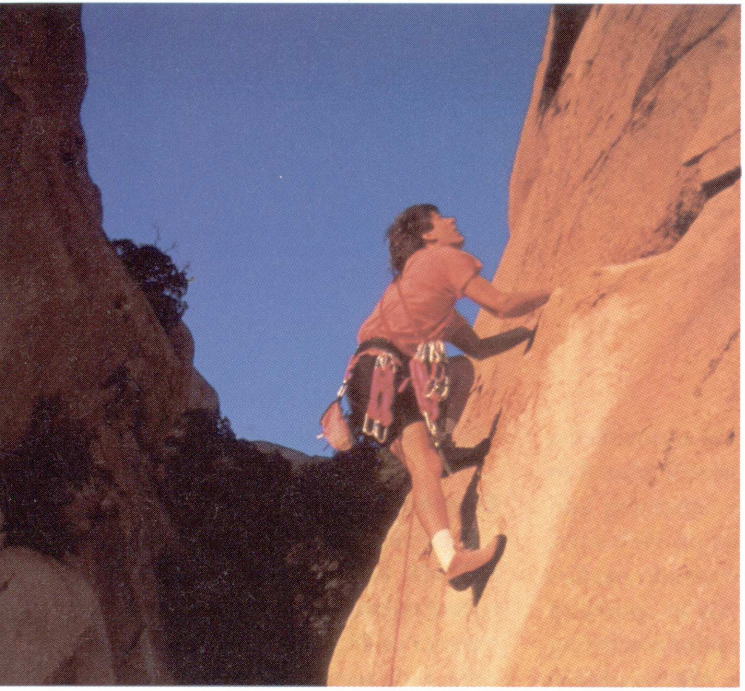
Shared Success for HP Computing Professionals

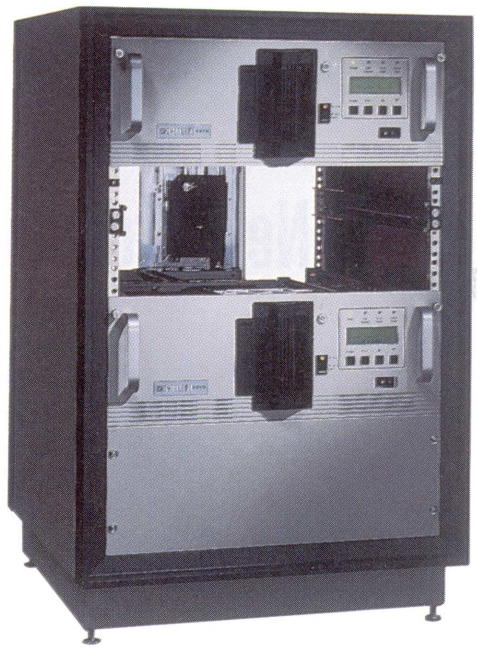
As a not-for-profit association of HP computing professionals, Interex is dedicated to meeting the information, education, and advocacy needs of its members worldwide.

In today's world of rapidly changing technology, Interex puts hands-on solutions to hardware, software, and operating system problems at your fingertips. Because members actively contribute—exchanging ideas and sharing solutions—Interex is a vital link in the transfer of HP expertise.

Operating independently from Hewlett-Packard, Interex has more than 20 years of serving HP computing professionals. Through its publications, conferences, and volunteer committee structures, Interex has the qualifications to represent you, a valuable member of the HP user community.

Interex® is a trademark registered in the U.S. Patent and trademark office.





Tough to scale

Easy to scale

Protect your investment in tape backup equipment with the LibraryXpress system. Begin with a base module containing one or two DLT4000 or DLT7000 drives and 10 cartridges. This will provide you with a storage capacity of up to 700 gigabytes and a data rate of 72 gigabytes per hour.



Scale up with any combination of stackable modules. Choose from a performance module, with up to five drives, or a capacity module, with 16 cartridges.

CIRCLE 33 ON READER SERVICE CARD

The exclusive Global Control unit passes cartridges between modules, allowing any cartridge to access any drive in the system, for true maximum system performance. Up to nine modules can be combined to create the ultimate automated storage system. Network administrators can continually fine-tune LibraryXpress expansion by balancing the ratio of tape drives to cartridges. Optimal configurations based on individual capacity, performance and budget requirements are easily achieved without the need to lock into a vendor-defined growth path of only a few possible configurations.

IEM is a full solution provider:
Combine LibraryXpress with our excellent software backup choices Alexandria and DallasTools.

IEM: Providing Solutions for a Lights-out Environment

In the U.S. and Canada:
IEM, Inc., P.O. Box 1889
Fort Collins, CO 80522 USA
Phone: (970) 221-3005
(800) 321-4671
Fax: (970) 221-1909

In the United Kingdom:
IEM, Inc., Unit 6, Salisbury House,
Wheatfield Way, Hinckley Fields,
Hinckley, Leicester LE10 1YG
Phone: +44 (0)1455 239000
Fax: +44 (0)1455 239668

All others:
IEM International Sales
1629 Blue Spruce Drive
Fort Collins, CO 80524 USA
Phone: +[1] 970-221-3005
Fax: +[1] 970-221-1909

© 1997 Anis Inc. LX1.0497

email: info@iem.com

http://www.iem.com



New Products

Web Site Tracking

Andromedia, Inc. has announced a partnership with Hewlett-Packard to provide solutions for reporting and understanding activity on Web sites. ARIA Recorder Reporter, Andromedia's Web site tracking tool, will be optimized for HP-UX through a joint marketing and technology alliance. ARIA combines capture of Web site activity data as it happens with ease of installation and maintenance. Automatically generated reports are available through an ordinary Web browser. ARIA sites may customize the information displayed to match their needs.

Contact Andromedia, phone: (415) 278-0700, http://www.andromedia.com.

Secure Web Server

C2Net Software, Inc. has announced Stronghold 2.0, a major upgrade to its secure Web server based on Apache. Stronghold now includes uncompromised security, Web-based configuration, and new protocol support.

Current "export" versions of Netscape and Microsoft Web servers use a weak 40-bit cipher. By comparison, all C2Net software can use at least 128-bit keys.

Stronghold is now fully compliant with the new HTTP/1.1 standard. HTTP/1.1 contains many new features, including improved content and language negotiation, improved persistent connections, and better recovery from interrupted transfers.

Stronghold 2.0 also includes a Web-based configuration manager, allowing administrators to administer sites from their chosen Web browser.

The new release includes support for Secure Sockets Layer (SSL) Version 3, which provides stronger security and more flexibility in choosing ciphers.

Contact C2Net Software, phone: (510) 986-8770, http://www.c2.org.

Empress RDBMS

Empress has announced Empress RDBMS V6.10. To increase speed, reduce memory usage, and optimize the transfer of BLOB data across the network, bulk segmentation facilitates the retrieval and update of specified portions of BLOB data in a record. For dynamic Web information and to augment the power of an HTML application, the Empress Report Writer can be used as a CGI program. A JDBC/ODBC Bridge Interface enables Java programmers to access Empress databases through the use of standard JDBC calls. For the real-time and near-real-time developers, Empress engineers have developed a microsecond timestamp that marks time to one millionth of a second and keeps dates to the year 10000.

Empress RDBMS V6.10 is available as a developer's toolkit, which also includes an HTML toolkit, ODBC Server and Client, and Dynamic SQL. Pricing for the Empress RDBMS V6.10 developer's toolkit starts at $1,000 for two users on Windows NT, $5,600 for typical workstations, and $22,400 for midrange servers. The Empress RDBMS runs on Windows and all major UNIX platforms, including HP-UX.

Contact Empress, phone: (301) 220-1919, http://www.empress.com.

ODBC Driver for COBOL

Acucobol, Inc. has announced AcuODBC, the Vision ODBC Driver. This new interface links COBOL data with Microsoft Windows graphics.

AcuODBC allows reads and updates to the vision data and provides flexibility for both end users and developers.

Applications need to be compiled only once to instantly plug into any ODBC-compliant data source—no restructuring or reengineering is needed.

AcuODBC operates in single-tier (stand-alone) and multitier (client-server or networked) environments; it is available in both 16-bit (for Windows 3.x) and 32-bit (for Windows NT and 95) versions; and it supports SQL commands and options by translating the ODBC SQL calls into Vision I/O commands.

Pricing for AcuODBC starts at $300.

Contact Acucobol, phone: (800) COBOL85 or (619) 689-7220, fax: (619) 566-3071, http://www.acucobol.com.

Database Backup Software

IEM, Inc. has announced DallasTools, a suite of layered products designed to facilitate the complete backup, restore, and archiving of a heterogeneous network environment using standard communication protocols and a wide range of storage devices. The product is able to back up data to any device connected to any UNIX server on the network and control it centrally from a DallasTools Cluster server.

DallasTools supports a wide range of storage devices, including all 8-mm tape drives (except 8200), all 4-mm DAT drives, 3480/90 style tape drives, DLT drives, and optical devices.

Contact IEM, phone: (970) 221-3005, fax: (970) 221-1909 (U.S.); phone: (+44) 01455 239000, fax: (+44) 01455 239668 (U.K.).

Embedded Web Server

Real Time Integration, Inc. has added an embedded Web Server inside the NetAcquire distributed data acquisition system. NetAcquire is an intelligent data acquisition system that makes analog and digital data available over

YOUR CONCERN FOR YOUR DATA INTEGRITY IS OUR BUSINESS

Proven Storage Solutions For Mission Critical Applications

WHEN DATA IS CRITICAL TO YOUR FUTURE, CHOOSE CONLEY'S HIGH AVAILABILITY SOLUTIONS

- Maximize your uptime
- Ensure uninterrupted access to on-line data
- Achieve maximum protection against data loss
- Easily monitor your system's status
- Competent emergency help available 24-7
- Guarantee system redundancy to eliminate crashes

If you lose your job, it won't be because you lost your data

Top financial institutions identify CONLEY as their premier choice for high availability RAID solutions

CONLEY

Conley Corporation 420 Lexington Avenue New York, New York 10170
Tel: 212.682.0162 Fax: 212.682.0071 email info@conley.com www.conley.com

CIRCLE 31 ON READER SERVICE CARD

SMART SOFTWARE FOR DUMB TERMINALS



FineLine Message Manager™

Full-featured electronic mail & appointment scheduling software for character-based UNIX systems. System mail & Internet compatible.

Pricing starts at only $200 for 5 users.

FineLine UnixBall™

Multiuser character-based action game for UNIX. Enables playing multiple terminals against each other in real-time.

Pricing starts at only $50 for 2 users.

FineLine Menu Manager™

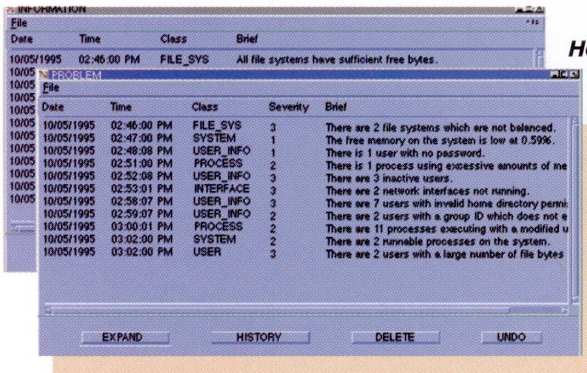
Professional menu generation software for creating character-based user menus which can execute any command with a single keystroke.

Pricing starts at only $75 for 5 users.

Download a free trial copy from the Web or call us for a 30-day evaluation copy complete with all documentation at no charge.

(800) 633-5250
(717) 688-9511
Web: www.soundideas.com
E-mail: info@soundideas.com

CIRCLE 168 ON READER SERVICE CARD



Heroix RoboMon Version 6.2

NT System Administration

Heroix Corporation has announced RoboMon Version 6.2 for Windows NT. The RoboMon NT software automates Windows NT system administration by monitoring and solving many problems associated with Microsoft's Windows NT Server, Exchange Server, Internet Information Server, SQL Server, and Systems Management Server.

RoboMon NT's Rule Wizard provides a fully integrated, graphical environment where users can easily customize RoboMon to monitor and solve site-specific problems. RoboMon monitors Windows NT based on rules that determine what conditions to look for and what notification or corrective actions to take when it encounters problems. It places schedules, selections, conditions, and actions at the user's fingertips in an easy-to-use, graphical form.

An enhanced set of out-of-the-box rules now includes new rules for checking directories, application failures, and changes to security policy.

RoboMon's e-mail notification action now supports Lotus' CC: Mail and SMTP mail as well as Microsoft Exchange. RoboMon rules can now page personnel through an integrated paging action. The SNMP Trap action has been enhanced to allow severity and status information to be sent to Network Managers such as OpenView.

Contact Heroix, phone: (800) 229-6500, ext. 232, http://www.RoboMon.com.

an Ethernet network. The NetAcquire embedded Web Server option allows for configuration and monitoring functions of hardware test and measurement applications from anywhere on a network. A Netscape Web client can easily administer all NetAcquire nodes on an intranet or the Internet. Configuration options include TCP/IP network settings, number and type of data channels, sample rates, refresh interval, and advanced server-side data processing.

NetAcquire embedded Web Server appears on a network as both a Web and an FTP server. Access control is also available using configuration password protection.

The NetAcquire Web Server is compatible with most Web browsers, including Netscape 3.0 and Internet Explorer 3.0. It is included with new NetAcquire hardware, which is priced at $5,495. Current users may upgrade for $495.

Contact Real Time Integration, phone: (888) 675-1122 or (425) 462-5817, fax: (425) 883-0463, e-mail: info@realtimeint.com.

NFS Client

WRQ has Reflection NFS Connection for Windows NT Version 6.1, which provides an integrated Windows NT and UNIX solution. It offers simplicity of use for non-technical staff and powerful functionality for expert UNIX users.

Windows NT users can now access file and print resources on UNIX networks, without knowing a single UNIX command.

Serious UNIX users can take advantage of features such as "r" commands and APIs for both NFS and RPC used for development of custom Windows applications.

WRQ's adaptive NFS optimizations automatically adjust to fluctuations in the network to maximize responsiveness throughout the day, making NT-UNIX connections efficient and highly stable.

Reflection NFS Connection for Windows NT is priced at $219.

Contact WRQ, phone: (800) 872-2829 or (206) 217-7100, fax: (206) 217-0293, e-mail: info@wrq.com.

OpenView Programming Book

Prentice Hall PTR and Hewlett-Packard Professional Books have announced the publication of *Power Programming in HP OpenView: Developing CMIS Applications*, by Raymond Caruso and Mark Kean. This book is a guide to leveraging HP OpenView's power in telecommunications network management.

HP OpenView is a powerful network management platform, but many network management applications fail to leverage its full power. Written for a wide range of telecommunications applications developers and project managers, *Power Programming in HP OpenView* will help readers develop robust telecommunications network management applications that leverage all of OpenView's capabilities, especially its support for CMIP and CMISE.

Power Programming in HP OpenView, by Raymond Caruso and Mark Kean, Hewlett-Packard Professional Books, 1997 (ISBN: 0-13-443011-5), is priced at $45.

Contact Prentice Hall PTR, phone: (201) 236-7139, fax: (201) 236-7131.

Data Warehouse Retrieval

Dynamic Information Systems Corporation (DISC) has announced OMNIDEX for Star Schemas, which provides performance improvements for sites

with large data warehouses that utilize a "star-schema" methodology. By installing the OMNIDEX multidimensional indexes on top of existing the star schema, end users can achieve enhanced retrieval performance.

The company states that OMNIDEX for Star Schemas has reduced ad hoc queries from eight hours to less than a minute against a star-schema "fact" table containing 80 million rows. OMNIDEX for Star Schemas integrates seamlessly between star schemas stored in popular RDBMs and popular analysis/reporting tools such as MicroStrategy's DSS Agent, Business Objects, and Cognos Impromptu.

OMNIDEX enhances multiple data structures including Oracle, Sybase, Informix, Rdb, RMS, C-ISAM, IMAGE/SQL, and flat files, and runs on multiple platforms, including HP-UX and MPE.

OMNIDEX is priced per CPU, from $40,000 to $150,000 per server.

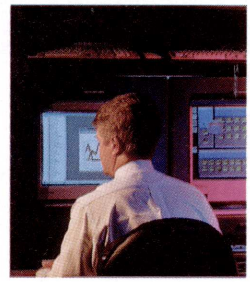
Contact DISC, phone: (303) 444-4000, http://www.disc.com/index.html.

Java Viewer

Hummingbird Communications Ltd. has announced Common Ground Java Viewer for DigitalPaper documents. Universally readable DigitalPaper electronic documents are printed from any Windows or Macintosh application and maintain their original appearance on the Web, in e-mail, fax, or hardcopy. The Java Viewer eliminates the need for manually downloading and installing a viewer application. Instead, DigitalPaper documents are now readable by anyone with a Java-enabled browser without any effort at all.

Contact Hummingbird Communications, phone: (415) 917-7300, fax: (415) 917-7310, http://www.hummingbird.com.

"SAP believes that PATROL® currently offers one of the leading solutions for in-depth monitoring and management of R/3. It is important that our customers have management tools like PATROL that can ensure optimal performance and high availability."



PATROL

Dr. Arnold Niedermaier,
Technology Marketing
SAP AG

Top industry leaders are relying on PATROL for a reason.
To find out why visit us on the Web at: www.bmc.com/patrol
Or call today: 800 811-6766

BMC® SOFTWARE
www.bmc.com/patrol

BMC Software, the BMC Software logos and all other product or service names are registered trademarks or trademarks of BMC Software, Inc. in the USA and in other select countries. ® and ™ indicate USA registration or USA trademark. Other logos and product/trade names mentioned are registered trademarks or trademarks of their respective companies. ©1997, BMC Software, Inc. All rights reserved.

CIRCLE 145 ON READER SERVICE CARD

Structured Application Development

Cayenne Software, Inc. has announced Version 7 of TeamWork, an analysis and design solution for structured application development. Over 3000 new enhancements enable developers to build technical, embedded applications with greater speed and efficiency.

New features of TeamWork Version 7 include: improved editing of graphical objects; support for industry-standard text editors; WYSIWYG printing for increasing the readability and value of Mosaic printouts; streamlined and flexible system administration; integration with Inso Corp's DynaText to facilitate Web browser viewing of online documentation and help; ADA design-to-code consistency checking; custom versioning that facilitates management of large databases; and automatic retrieval of locally locked objects.

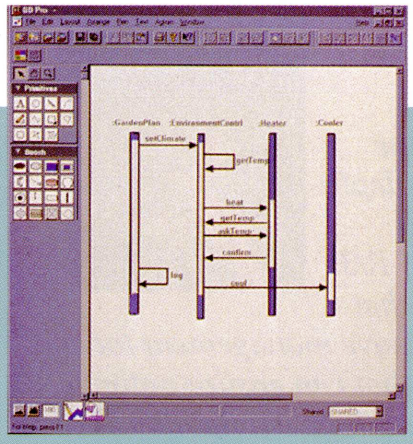
TeamWork is an integrated set of tools and services that enable software developers to create, store, review, and maintain structured specifications and data models. The product supports most UNIX platforms, including HP-UX.

Pricing for TeamWork Version 7 starts at $9,600.

Contact Cayenne Software, phone: (617) 505-3595, fax: (617) 229-9904, http://www.cayennesoft.com.

Client-Server Management Software

Power Center Software has announced its "Test Drive" Program for its Power Center suite. The program allows system and network managers to "test drive" the product on their systems for 30 days and entitles them to a significantly lower entry price for the entire Power Center suite if ordered after the "Test Drive." Through the "Test Drive" Program, four value-added products are included when the System Problem Monitors are licensed.



ASTI GDPro

Software Design

Advanced Software Technologies, Inc. (ASTI) has announced GD*Pro*, a cross-platform development tool that supports various software design methodologies and workgroup collaboration.

Developers can create object-oriented software designs. GD*Pro* supports collaborative multiuser development—groups of geographically separated users can run GD*Pro* over an internal corporate network or the Internet and interact "live" with a common software design.

GD*Pro* generates source code for all supported design methodologies in one or more languages, including C++, C, and Java. With GD*Pro*, Java, C++, and C source code can be reengineered to automatically produce detailed design diagrams. GD*Pro's* reengineering support is intended to help developers redesign existing software, convert existing software from one language to another, and to support re-hosting and client-server development.

GD*Pro* supports PCs running Windows 95 and Windows NT and most UNIX workstations.

Pricing starts at $995 for either Windows version and $3,495 for the UNIX version.

Contact Advanced Software Technologies, phone: (303) 730-7981, fax: (303) 730-7983, e-mail: info@advancedsw.com.

"Test Drive" offers significant fault and availability monitoring, as well as automated correction and notification in order to provide consistent service availability.

Power Center provides enterprise-wide automation and management of multivendor client-server enterprises.

Contact Power Center, phone: (303) 220-9789, fax: (303) 220-9776, http://www.powercenter.com.

New from BASIS

ODBC Driver

BASIS International Ltd. has announced Release 1.1 of its ODBC Driver, which can unlock BB Business BASIC data files previously inaccessible to Windows users.

The driver can now read and write non-normalized as well as normalized data files. (Non-normalized files are ones where fields may have subfields, data may be aligned on half-byte boundaries, and more than one type of record may exist in the same file.)

It has been a severe restriction that until now, software developers and data analysts could not take advantage of the full power of ODBC-standard database access and ODBC-compatible applications unless data files were first "normalized."

The new driver works in mixed platform client-server environments.

It is licensed for $150 per copy.

Business Basic Enhancements

BASIS International has announced that more than 100 improvements, tools, and new features have been added to BASIS' Visual PRO/5 Business Basic software for Windows, UNIX, and PRO/5 Data Server. Four significant advancements are: Program Wizard, Print Preview, ResBuilder, and DDBuilder.

Program Wizard is an application development tool that produces programs for file maintenance or report generation. Print Preview eases the burden of developing complex Windows code. ResBuilder supports the rapid development of Windows resource files in BRF format. DDBuilder is an advancement of the Data Dictionary.

Prices for a Business Basic development system based on Volcano begin at approximately $1,500 for a 10-user license.

Contact BASIS International, phone: (505) 345-5232, fax: (505) 345-5082, e-mail: info@basis.com, http://www.basis.com/.

Object Repository

Softlab Enabling Technologies has announced Enabler, which changes the repository from storage device to information management device in the Windows 95/NT and UNIX enterprise. The Enabler repository's open architecture enables it to act as a hub or supergroup, linking different tools and disparate applications into a single, common information management infrastructure. Users can share and re-use product and application components throughout their organization.

Enabler allows tools, application, and product vendors to provide a solution that scales to the enterprise. It will enable more effective component reuse and management throughout the enterprise.

Enabler is priced at $3,500 per seat.

Contact Softlab, phone: (770) 668-8811 (U.S.), (+49) 89 99360 (Germany), e-mail: info@softlabna.com, http://www.softlabna.com.

New from Bradmark

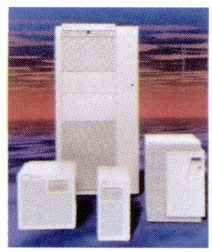
Graphical Display Tool

Bradmark has announced DBGENERAL Server Manager for Oracle, a graphical display tool that shows the logical and physical structures of an Oracle tablespace and predicts space problems with trend analysis. DBGENERAL is a comprehensive toolset for monitoring and tuning Oracle database servers across distributed enterprise networks.

Features include tablespace management, trending analysis, user and role management, structural change management, database administration, and SQL generation.

The product's script file generation facility allows the user to execute any changes in batch or online. It also can

"HP



is reselling PATROL® technology and products because of the tight integration BMC Software is delivering through the HP OpenView enterprise solutions. Together, HP OpenView and PATROL pro ide the highly desired complementary functionality that our customers need today in the areas of application and data management."

PATROL

Johnnie-Mike Irving, General Manager
HP Professional Services Organization
for the Americas

Top industry leaders are relying on PATROL for a reason.
To find out why visit us on the Web at: www.bmc.com/patrol
Or call today: 800 811-6766

BMC® SOFTWARE
www.bmc.com/patrol

BMC Software, the BMC Software logos and all other product or service names are registered trademarks or trademarks of BMC Software, Inc. in the USA and in other select countries. ® and ™ indicate USA registration or USA trademark. Other logos and product/trade names mentioned are registered trademarks or trademarks of their respective companies. ©1997, BMC Software, Inc. All rights reserved.

CIRCLE 146 ON READER SERVICE CARD

Business Mapping Applications

MapInfo Corporation has announced integrated mapping products that bring geographic analysis into corporate information systems. Application developers and database analysts can use mapping technology with corporate databases such as Oracle, Informix, and Lotus Notes applications. These maps can reveal trends and relationships to support critical decisions.

MapX, an OCX mapping component, lets developers embed mapping into common business applications. It is a 32-bit program, runs on Windows NT and 95, and is integrated using Visual Basic, Visual C++, Delphi, Lotus Script, or PowerBuilder.

SpatialWare is server software that works with database management systems to store geographic data in relational databases along with powerful query functionality. It is available on HP-UX and other UNIX operating systems.

MapMarker is real-time, intelligent address-matching software with enterprise-wide scalability. It is available for Windows 95 and NT.

MapX's pricing starts at $4,960; SpatialWare's pricing starts at $20,000; and MapMarker's pricing starts at $2,595.

Contact MapInfo, phone: (518) 285-6000, fax: (518) 285-6060.

MapInfo Professional



Data Exchange Manager

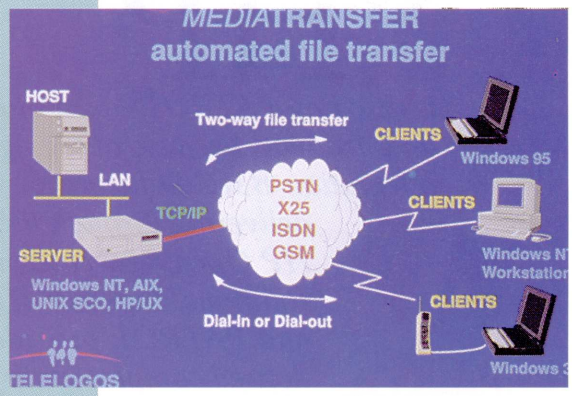
TELELOGOS has announced Mediatransfer, which automatically manages data exchange between a central server and remote sites. Two versions are available: the Station version takes calls one at a time (for isolated machines on the network), and the Server version manages up to 64 calls at a time through one or more networks (for computers installed as communications gateways).

After users plug into the network, mobile telephone, or wall phone jack, systems are connected and files are sent and received automatically without further user intervention. Data can be from word processors, spreadsheets, e-mail, database management packages, or other application software. All types of data are sent in a single exchange.

Transmission speeds range from 2,400 to 128,000 bps. Data is sent in batches to increase transmission speed.

Mediatransfer is available on DOS, Windows, and UNIX. Data is exchanged through X25, TCP/IP, and ETP protocols. The company is seeking distributors for the North American market.

Contact the French Technology Press Office, phone: (312) 222-1235, fax: (312) 222-1237 (U.S.).



TELELOGOS
Mediatransfer

extract the DDL from the ORACLE catalogue for any object and has an SQL editor. Other features are: user, role, and authority management; tablespace management and table reorganization; structural change management; general database administration; and SQL script generation.

The Server Manager V2.0 for Oracle is available for Windows NT, 95, HP-UX, and others.

A 30-day, free evaluation demonstration is available by calling (800) 275-2723.

Performance Monitor for Sybase

Bradmark has also announced its DBGENERAL Performance Monitor for Sybase, which enables management of geographically dispersed, complex database environments across corporate networks.

Administrators can monitor all Sybase SQL Servers within a network environment, regardless of server location. Using distributed agents running unattended on the server, the Performance Monitor continually analyzes the SQL servers, sending alerts to performance anomalies.

Key features include: monitoring of unlimited servers, from any clients; identification of problems instantly, using powerful drill-down technology; measurement of any performance metric with either predefined events or user-defined events; specification of multiple actions for notification of significant server events and the viewing of snapshot, trending, and repository data in a custom presentation format.

The Performance Monitor V2.0 for Sybase is available for Windows NT, 95, HP-UX, and others.

For a 30-day free evaluation demonstration, call (800) 275-2723.

Contact Bradmark, phone: (800) 621-2808, fax: (713) 621-1639, http://www.bradmark.com.

Backup and Recovery System

Workstation Solutions has announced Quick Restore V2.2, a backup and recovery system for heterogeneous Windows NT and UNIX networks. From a central management console, Quick Restore automates backups to a wide variety of storage devices. Restore Browser GUI displays a seamlessly integrated view of Windows NT and UNIX files.

Quick Restore's custom device drivers maximize performance from tape drives "direct-to-block positioning." During the backup process, the physical block address of each file on the backup tape is recorded. When the file is restored, the tape spins at maximum speed to the exact file position, resulting in the fastest possible restore performance.

Server software packages are priced at $5,400. Client software for UNIX is priced at $1,800 per type of client architecture.

Client software pricing for Windows NT is per client, ranging from $2,250 for five clients to $15,000 for 200 clients.

Contact Workstation Solutions, phone: (603) 880-0800, http://www.worksta.com.

New from Xi Graphics

Graphic Software

Xi Graphics has announced AcceleratedX Server Version 3.1, which makes available workstation-class features like multiple visuals, overlays, and gamma correction for industry-standard PC hardware.

Version 3.1 provides a solution to graphics-related problems experienced

by Linux and Free BSD users, such as hardware support, color correction, and performance. It is also the foundation for Xi Graphics' new generation of high-performance desktop products.

AcceleratedX Server Version 3.1 drives popular graphics hardware, supporting more than 450 different cards.

Text-based configuration files permit easy modifications. When loaded, the product seamlessly replaces the installed X server and is fully compatible with the native X Windows System. The server is CDE-ready, enabling successful installation and higher performance of the Common Desktop Environment.

Network Desktop Solution

Xi Graphics has also announced CDE Business Desktop, a complete networked operating system that enables systems administrators to integrate industry-standard PCs into existing UNIX networks with better internationalization than, and the same functionality as, more expensive workstations.

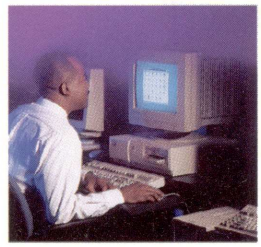
The Common Desktop Environment (CDE) provides users with a consistent GUI for accessing data and applications from anywhere in the network.

Xi Graphics' CDE Business Desktop includes the latest AcceleratedX Server,

Continued on Page 90

Intel

Corp. is integrating PATROL® management technology into the i960RP I/O processor in order to provide the direct management capabilities for I/O subsystems customers are asking for.



PATROL

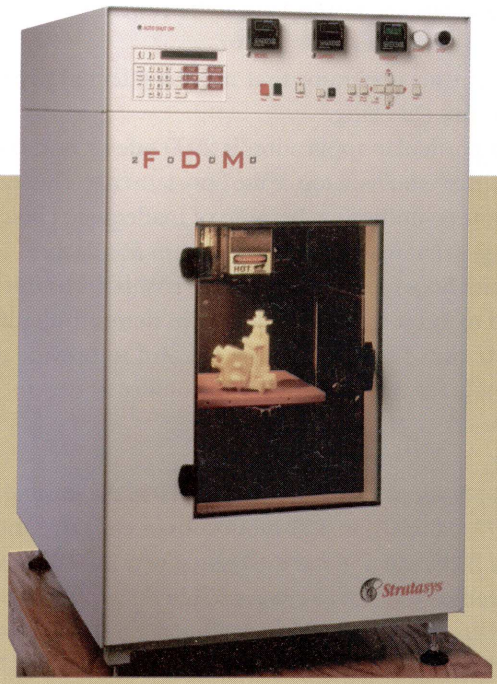
Top industry leaders are relying on PATROL for a reason.
To find out why visit us on the Web at: www.bmc.com/patrol
Or call today: 800 811-6766

BMC® SOFTWARE

www.bmc.com/patrol

BMC Software, the BMC Software logos and all other product or service names are registered trademarks or trademarks of BMC Software, Inc. in the USA and in other select countries. ® and ™ indicate USA registration or USA trademark. Other logos and product/trade names mentioned are registered trademarks or trademarks of their respective companies. ©1997, BMC Software, Inc. All rights reserved.

CIRCLE 147 ON READER SERVICE CARD



Stratasys FDM2000

Rapid Prototyping System

Stratasys, Inc. has announced the FDM2000 rapid prototyping systems, which allow for fast throughput. It also announced the automated QuickSlice 4.1, which provides users with the ability to control build parameters. This enables the FDM2000 to achieve enhanced part feature definition and surface finish.

QuickSlice 4.1's pre-defined parameter files, at a range of layer thicknesses from .178 to .356 mm, enable the system to create parts with enhanced feature detail. In addition, the FDM2000's head moves at a rate of up to 1.5 inches per second. The FDM2000's fast throughput is achieved through controller, head, and Z-stage design enhancements.

The system operates on HP, Silicon Graphics, Sun, and Windows NT workstations.

Contact Stratasys, phone: (612) 937-3000, fax: (612) 937-0070, e-mail: fdm@stratasys.com, http://www.stratasys.com.



Question & Answer

Q: What is being done in HP-UX for the Year 2000 problem?

A: Virtually every company will be affected by the year 2000 problem. For the most part, the issues are with legacy programs and data files that used only 2-digit or single byte values to represent the year, with an assumption that the current century is 1900. Unfortunately, these programs and/or the data files cannot handle the year 2000 and beyond. Ideally, all these programs will be scrapped and rewritten by that time.

However, some operating system limitations will affect even correctly written programs. HP-UX versions prior to 10.20 do not currently handle the year 2000 correctly. This includes all versions of 9.xx and earlier. It is important to note that simply setting the date to *Dec 31, 1999* and letting it roll over is a very inadequate test for any operating system. There are subtle pieces of code in the file system, utilities that may use dates, such as *awk* and *find*, opsystem services such as *cron* and networking, that all need careful examination.

The objective of the 2000 Rollover Program in 10.30 is to ensure HP-UX becomes Year 2000 Safe (works beyond the year 2000). The scope of this program is at the opsystem core level. It does *not* imply that all applications that are HP-UX-based will be year 2000 safe. Here are some of the reasons why:

It is clear at this point that during the 2000 Rollover Program, one of the major problems is the misinterpretation of the date field by HP-UX code at the core level. Thus, it is highly possible the same kind of error lies in the application-level code.

Applications may have their own defined "date" data field. These date fields, as well as all of their dependent applications, would have to be rewritten in order to avoid the year 2000 problem.

There are some library calls in user programs that may be affected by the updated code in *libc*, which will require checking.

So, in order to determine if a particular application is 2000 safe or not, the programmer needs to look at the code that is built on top of the opsystem calls. This includes date calculations that may overflow or underflow because of inadequate choices for numeric precision as well as changes that may have taken place in *libc* patches.

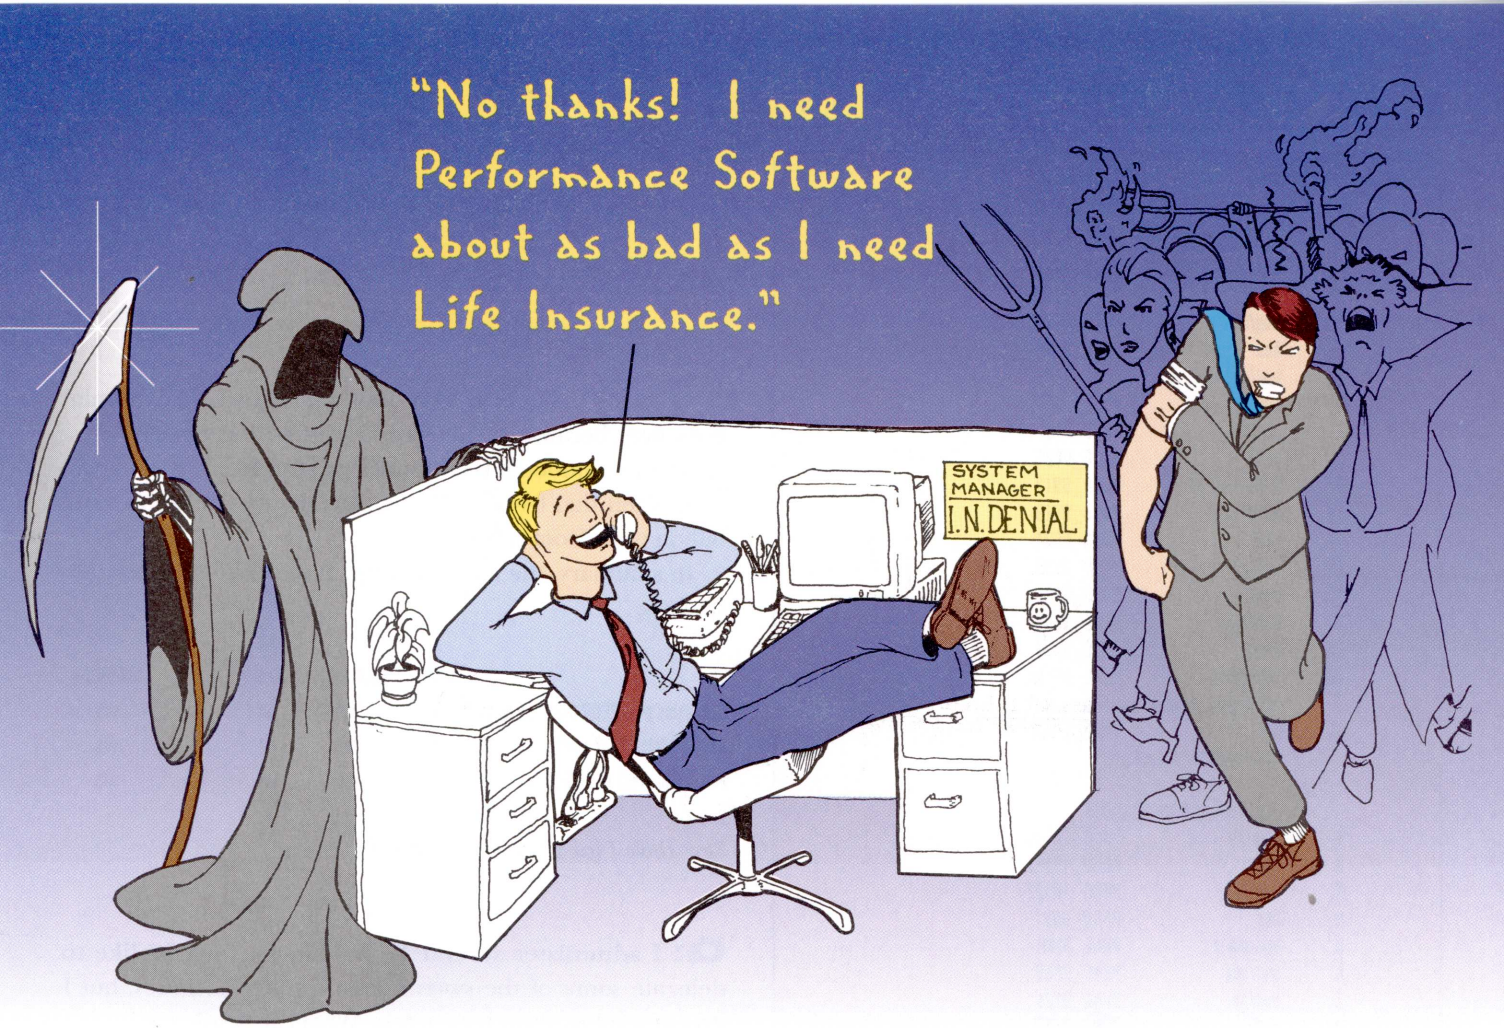
Programmers should obtain the *libc* white paper from HP's Web site as the patch ID: PHCO_10175. It is intended to aid software developers in determining if their software has a dependency upon any of the changes to *libc*, and to provide guidelines for making the software ready for the year 2000 with respect to *libc*.

How HP-UX Calculates Time

HP-UX converts dates from applications to seconds, with 1970 as the starting point (more specifically, January 1, 1970, 0:00:00). In a 32-bit OS, this means that the time counter is "maxed out" in the year 2038, with 2037 as the highest year HP-UX supports.

Presently, HP-UX 10.20 and earlier will not always calculate dates after 2000 correctly. This shortcoming will be corrected in HP-UX 10.30, due out in 1997. However, because 10.30 is still based on 32-bits internally, it is still subject to the 2037 time maximum. Keep in mind that 2037 is about 40 years away, a time frame where the design of future computers and opsystems can only be guessed.

Continued on Page 18



It's Time For Your Performance Review

System performance management demands a great deal. You must deal with current workloads while optimizing user response times. Simultaneously, you must plan for the future, anticipating where new performance problems will occur. **SOS/9000 Performance Advisor** for HP-UX and **SOS/X Performance Advisor** for X Windows environment will help you maximize your 9000's performance as well as provide you the information needed to effectively plan for the future.

- Take proactive control of your system's performance with very little effort.
- Better utilize system resources and improve efficiency of the entire data processing operation.
- Know the track record of your system's performance to determine future resource needs.
- Save thousands of dollars over other tools.
- Let **SOS/9000 Performance Advisor** be your system performance "watchdog."

For more information or a Free demonstration copy
Call Lund Performance Solutions Today! 541-926-3800

LUND
PERFORMANCE SOLUTIONS
Performance Beyond Expectation

240 2nd Ave. SW • Albany, OR 97321 USA
(541) 926-3800 • (541) 926-7723 (fax)
E-Mail: info@lund.com
Visit our Web Site: www.lund.com
CIRCLE 58 ON READER SERVICE CARD

SOS/9000
Performance Advisor

TABLE 1 *Year 2000 Patches*

	Year 2000 patches for 10.01	
1.	cron/at	PHCO_10112
2.	date	PHCO_10111
3.	rlog/co	PHCO_10110
4.	libc	PHCO_9419
5.	inet/mailx	PHNE_9379
6.	inet/elm	PHNE_9859
7.	inet/ftp	PHNE_9181
8.	netTL	PHNE_9424
9.	SSHA/mkboot	PHCO_9012
10.	SSHA/sar	PHCO_8819
11.	LVM	PHCO_8624
12.	audisp	PHCO_7933
13.	SAM	PHCO_8131
14.	Cold Inst	PHCO_9648
15.	kermit	PHCO_9497
16.	kernel/HIL	PHKL_10143
	(s700 only, no s800 patch required)	
17.	libc	PHCO_10175
	(Compatibility White Paper)	
	Year 2000 patches for 10.10	
1.	cron/at	PHCO_10120
2.	date	PHCO_10121
3.	rlog/co	PHCO_10122
4.	libc	PHCO_8981
5.	inet/mailx	PHNE_9379
6.	inet/elm	PHNE_9859
7.	inet/ftp	PHNE_9181
8.	netTL	PHNE_9425
9.	SSHA/mkboot	PHCO_9013
10.	SSHA/sar	PHCO_8819
11.	LVM	PHCO_8625
12.	audisp	PHCO_7934
13.	SAM	PHCO_8130
14.	Cold Inst	PHCO_9648
15.	kermit	PHCO_9507
16.	kernel/HIL	PHKL_10142
	(s700 only, no s800 patch required)	
17.	libc	PHCO_10175
	(Compatibility White Paper)	
	Year 2000 patches for 10.20	
1.	cron/at	PHCO_10123
2.	date	PHCO_10124
3.	rlog/co	PHCO_10125
4.	libc	PHCO_8979
5.	inet/mailx	PHNE_9381
6.	inet/elm	PHNE_9860
7.	inet/ftp	PHNE_9785
8.	netTL	PHNE_9081
9.	SSHA/mkboot	PHCO_9014
10.	SSHA/sar	PHCO_8820
11.	LVM	PHCO_8626
12.	SAM	PHCO_8133
13.	kermit	PHCO_9508
14.	kernel/HIL	PHKL_9921
	(s700 only, no s800 patch required)	
15.	kftpd	PHNE_9786
16.	libc	PHCO_10175
	(Compatibility White Paper)	

Once HP-UX goes to 64-bit, this maximum time constraint goes away, because the allowed time range increases by 2^{32}, a very large number. It is doubtful that any HP-UX systems will exist when that 64-bit value reaches the maximum (several billion years from now).

In summary, the current (unpatched) HP-UX opsystem status is:

	9.XX	10.01/.10	10.20	10.30	64-BIT+
Calculates post-2000 dates correctly	no	no	no	yes	yes
Handles dates 2037 and beyond	no	no	no	no	yes
Patches to meet Year-2000 compliance	no	yes	yes	n/a	n/a

See *Table 1* for a list of year 2000 patches.

Q: I administer several 10.X systems. I would like to delegate some of the system administrator activities, but I do not want to grant root access. What can I do?

A: SAM can be configured to provide a subset of its functionality to certain users or groups of users. It can also be used to build a template file for assigning SAM access restrictions on multiple systems. This is done through the Restricted SAM Builder.

You can access the Restricted SAM Builder by invoking SAM with the -r option. In the Builder, you may assign subsets of SAM functionality on a per user or per group basis. Once configured, you can use the -f option to verify that the appropriate SAM functional areas are available to the specified user.

A non-root user who has been given Restricted SAM privileges simply executes */usr/bin/sam* and will see only those areas that you configured that user to access. SAM will automatically give root privilege for those tasks that require it.

Q: We develop our own software packages. In the past, we have allowed all users to install and remove the software. We would now like to restrict access to the depots. How can this be done?

A: You can use the *swacl(1M)* command. This command allows you to modify the Access Control lists associated with software products. There are many levels of SD ACLs. *Table 2* is a list of commands for each level and the associated

TAMING UNIX

An Introductory Guide to Performance Management for the HP-UX System Administrator

By Robert A. Lund

A WISE INVESTMENT FOR ONLY $59.95



You, the humble and righteous system administrator, Whilist in the midst of performing your duties, are nearly overcome by a herd of renegade daemons.

Some of the Topics Covered Include:

- Hardware/Application Interplay
- Memory Management
- Load Balancing Techniques
- Capacity Planning
- HP-UX Kernal Tables
- Disk Space Management

"In this book, Bob Lund offers a common sense approach to the sometimes mysterious topic of UNIX system performance. Recommended reading for all system administrators. A great addition to any data processing professional's toolbox."

LEROY RUGGERIO
Informational Technology Engineer, Data Based Systems

CIRCLE 62 ON READER SERVICE CARD

LUND
PERFORMANCE SOLUTIONS
Performance Beyond Expectation

Call and order your copy today!

(541) 926-3800

240 Second Ave. SW • Albany, OR 97321 • Phone (541) 926-3800 • Fax (541) 926-7723
e-mail: info@lund.com • Visit or web site at: www.lund.com

TABLE 2 *swacl Command*

`swacl -l global_product_template`	
/var/adm/sw/security/_PROD_DFLT_ACL :	default depot product template used to initialize the product acl template for a new depot
`swacl -l global_soc_template`	
/var/adm/sw/security/_SOC_DFLT_ACL :	acl template for new depot
`swacl -l host`	
/var/adm/sw/security/_ACL :	host acl info
/var/adm/sw/security/_OWNER :	owner info for the host object
`swacl -l depot @ <depot>`	
/<depot>/catalog/dfiles/_ACL :	depot acl
/<depot>/catalog/dfiles/_OWNER :	owner info for the depot object
`swacl -l product_template @ <depot>`	
/<depot>/catalog/dfiles/_PROD_DFLT_ACL :	template for new products created in <depot>
`swacl -l product <product> @ <depot>`	
/<depot>/catalog/<product>/pfiles/_ACL :	product acl
/<depot>/catalog/<product>/pfiles/_OWNER :	owner info for the product object

storage files. These commands display the existing ACL information. Please see the *swacl(1M)* command for information regarding modification.

Q: What do the parameters *maxssiz*, *maxtsiz*, and *maxdsiz* represent?

A: *maxdsiz* is the most commonly modified of the three. It represents the maximum data storage segment size for an executing process. This segment contains fixed data storage such as globals, arrays, space allocated using *malloc()*, and such. Whenever the system loads a process or an executing process attempts to expand its static storage segment, the system checks the size of the process's data storage segment. If the process's requirements exceed *maxdsiz*, the system returns an error to the calling process. So, increasing *maxdsiz* does not cause an increase in memory usage; it allows for an increase.

maxssiz represents the maximum size of the dynamic storage segment, also called the user-stack segment, or an executing process' run-time stack. This segment contains stack and register storage space. This parameter does not usually need to be increased.

The stack grows dynamically. As it grows, the system checks the size of the process' stack segment. If the stack size requirement exceeds *maxssiz*, the system terminates the process.

maxtsiz defines the maximum size of the shared text segment (program storage space) of an executing process. Program executable object code is stored as read-only, and thus can be shared by multiple processes. It is very seldom that this parameter needs to be increased.

Q: After updating to 10.20, sendmail fails to work properly. Is there some additional configuration that is needed? I am not using DNS.

A: Check the following on your system:

1. Make sure that */etc/mail/service.switch* has the proper entry for these files:
 hosts files
2. Make sure that */etc/mail/sendmail.cf* has the *Dj$w* macro set to the qualified system name.
3. Make sure that *etc/mail/sendmail.cf* has the *O* option for the service.switch file uncommented:

```
O ServiceSwitchFile=/etc/mail/service.switch
```

4. Make sure that the first line of */etc/hosts* has the fully qualified name of the system.
5. Add '*<hostname>*.' as an alias for the local host in */etc/hosts*; don't forget the final dot.
6. Make sure that */etc/mail/sendmail.cw* has three lines:

```
short host
fully.qualified
local host
```

7. Kill and restart sendmail:

```
# sendmail -bd -q30m
```

Q: I would like to create a network install server for my new Series 700 workstations. What are the steps to do so?

A: Follow the these guidelines:

- The server can be a 700 or 800 system. However, it must be running the same level OS as that you are installing.
- The depot server can be the same system as the install server or a different system. It must be running 10.X.
- There must be one install server per subnet.
- The HP-UX Install product or the NetInstall bundle must be loaded on the server.
- The HP-UX Install product requires about 21 MB of free space in the server's */usr/lib* directory.
- Edit the */etc/instl_boottab* on the server and add at least one IP address that is reserved for booting install clients. The IP addresses you add should be for cold installs only and should not be used by any other systems. However, if you know the LAN Link Address (LLA) of the systems you will install, you can use the IP addresses of these systems if you use the reserved keyword.
- If you want to boot multiple systems from the server at once, you must add more than one IP address to this file. The server may deny boot services if multiple systems try to use the same IP address during booting.
- There are comments in the */etc/instl_boottab* file to help you edit this file. Also, see the *instl_bootd(1m)* man page.
- If you have problems booting the systems, review */var/adm/syslog/syslog.log* for error messages that indicate the type of problem.
- Refer to the following manuals for more information:
 Installing HP-UX 10.01 (10.10 or 10.20)
 Managing HP-UX Software with SD-UX

Q: I have a 735 workstation running HP-UX 10.20. I have a very large directory with a great number of small files. I found that I was unable to copy the contents of the directory to another via the *cp(1)* command. I received the following error:

```
# cp ../big/* .
sh: /usr/bin/cp: The parameter list is too long
```

How can I copy these files?

A: It is most likely that your problem is caused by the finite length of any executable's argument list. Previously, the maximum total argument list size was 20 kilobytes. It is important to remember that the environment also constitutes the argument list.

One immediate way to remedy the problem is to reduce the number of arguments that *exec(2)* must process. For instance:

```
# find ../big -exec cp {} . \;
```

Recently, a patch was released to address this issue, PHKL_10176. It has since been superseded by PHKL_10689. This patch allows for the possible allocation of 100 20-KB buffers to produce a maximum of 2,048,000 bytes for argument/environment information. ◧

General HP-UX questions are answered by Bill Hassell, a support engineer at the HP Atlanta Response Center. He can be contacted via e-mail at blh@hpuerca.atl.hp.com. Workstation questions are answered by Susan Potter, an HP-UX system support engineer in the Atlanta Response Center. Her e-mail address is sup@atl.hp.com.

HP-UX System Administration

by Chris Curtin

Root Password

THIS MONTH I AM GOING to present a solution to providing the root password to users when the administration staff is not available.

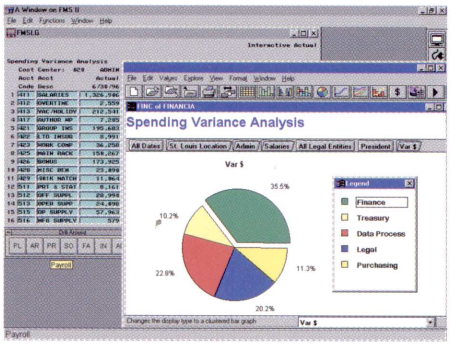
As the editor indicated a couple of months ago, I was unable to write one of these columns because I was out of town working on a large project. (No, system administration is not my primary job.) I spent three months on a factory automation start-up. During that time I was able to dial into my main server two or three times a day to check e-mail and make sure the systems were running. However, I was not able to be reached quickly by my colleagues back in Atlanta if they needed system administration help, since I was usually on the factory floor in a hard-hat and a hair net. (Yes, it was pretty humorous, but everyone else looked funny too.)

While many of you probably don't leave the office for such a long period of time, or travel where you can't be reached easily, what about when you are on vacation (or sitting in traffic for hours on end)? What if someone needs to get into your systems as root to address some critical problem? If you are like me, security is the major day-to-day concern. (If it isn't, maybe your major problems are caused by too many people being able to modify your system?) I was also apprehensive giving the root password to anyone, even if they had a good reason for having it.

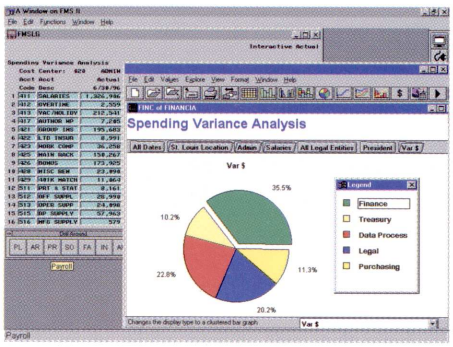
I hadn't thought much about my trip until a recent thread in the USENET *comp.unix.admin* newsgroup. Someone asked the question, How do I give root access to users in an emergency? (I can't find the initial post, but that was basically the question.) What the poster was looking for was suggestions for placing the root password on paper, somewhere that it could be accessed, but in such a way that he would know when someone accessed it.

After reading the first couple of posts in this thread, I wondered what would happen if neither I nor the others who knew the password could be reached. I knew I had some "back doors" built into our system which allowed execution of commands as root by non-root users. But the list of commands was limited and the back doors wouldn't allow the user to log in as root. Also, many commands were cumbersome to execute.

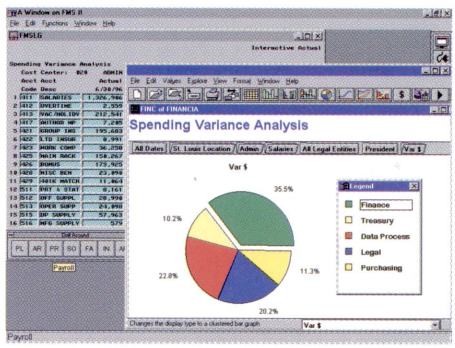
To address this problem, I wrote the root password on a piece of paper, put it in a security envelope (the kind with patterns on the inside that prevent you seeing the contents if you hold the envelope up to a light), signed my name across the flap, and taped it shut. Nowhere on the envelope did I write "root password" or anything else that would indicate what was in it. What I did write on it was a quick description, in terms a local person would understand, of what it was. For example: "The key to Elvis' bedroom" or "The honey for boo-boo." In both cases, Elvis and boo-boo are computers, but an outsider wouldn't know that. Next, I gave it to the office manager to place in the on-site storage safe. The only people who had access to this safe were the owners of the company, the office manager, and the receptionist. (You need to know the receptionist to understand why she had a key.) The basic rule was: If someone needs the root password, and one of the administration staff is not around, give him the envelope. Leave voice mail for me and the other staff and get the envelope back when the user is done. The user was supposed to send me e-mail about why he needed access.



HP 3000



HP 9000



Windows NT

Do you see a difference? Neither do our users!

***FMS II, OUR CLIENT/SERVER ACCOUNTING SYSTEM**, is portable across a variety of operating environments, including the HP 3000, HP 9000 and Windows NT.*

Even better, FMS II's Windows user interface looks exactly the same in all operating environments. So if you change your hardware, you don't have to change your software. In fact, you don't even have to retrain your users.

The bottom line is that FMS II has a lower long-term cost of ownership, which makes your bottom line look even better. For client/server accounting, there's simply no better solution. Call us at (800) 237-0028 to find out how we can help you.

MITCHELL HUMPHREY & CO.

EXCEPTIONAL SOFTWARE FOR EXCEPTIONAL ORGANIZATIONS

11720 Borman Drive
St. Louis, MO 63146
(800) 237-0028 or (314) 991-2440 Fax: (314) 991-5288
web address: www.mhco-fms.com
e-mail: mhmarket@mhco-fms.com

HEWLETT PACKARD Premier Solution Partner

Microsoft SOLUTION PROVIDER

CIRCLE 38 ON READER SERVICE CARD

This worked great, except for one thing: What if none of the people with a key was available? Rare yes, but it did happen. I never did address this issue before I left a couple of months later.

Before I present some other suggestions, I must point out that a user doesn't need the root password to have unlimited access to a system if he has physical access to the system. With all but the most secure versions of UNIX, a user with physical access to the server can reboot it and gain single-user access while the system is booting. This is why a physically locked computer room is so important to a strong security plan.

First some ideas for providing root-level access to users who need it for specific tasks, for example backups, mounting CD-ROM drives, rebooting the system:

- Obtain *sudo* from one of the various HP-UX public domain software sites. This program allows you to assign specific users specific commands they can execute as root. It also provides a nice audit feature that sends the real root user e-mail each time it is used (or in a batch every couple of hours, which is more typical). It is provided in source code form, so you can look it over before loading it on your system.
- Use the SAM functionality in HP-UX 10.X to provide SAM access to non-root users. While I have never used it (and generally don't use SAM for day-to-day tasks), I have heard good things about its use for backups by non-root users.
- Write your own *sudo*-like program. While this may have been an option a few years ago (before *sudo*), I strongly recommend against it. You have the *sudo* source code, so make sure you understand what it does and remove any features you don't like. Also remember you don't have to call the command "*sudo*."

Next, some ideas, some of which are odd but might work, and some of which are proven to work, for providing root access any time to non-privileged users:

- Modify *sudo* to allow any command by any user. Leave the auditing capabilities so you can track who did what. There are huge security issues with this approach, but it works.
- Implement a scheme similar to the one I used, putting passwords in an envelope. Again, the issue I described above would need to be addressed.
- Implement the envelope scheme, but place the envelope in a common location (e.g., on page 300 of some programming book). This has the downside of requiring the administration staff to check that location frequently to see if the envelope has been opened, and you have no history of who opened the envelope.
- Implement the envelope scheme, but place it on a non-system console terminal, where the administrator, or group of users would notice immediately if the envelope was missing. Same problems as previous idea.
- Provide all administration staff with 24-hour beepers. If you do this, make sure you pay them for carrying the beeper, a per-day rate and a per-call rate, and implement a disciplinary policy for punishing users who beep them if it is later determined that the administrator was not needed to address the issue. (I've heard too many horror stories of angry operators paging the system administrator at all hours because of some personal issue.)
- Implement a lock-box scheme which requires a key to access. While everyone would have the key, you would still have the audit trail issues.
- Implement a lock-box scheme that requires two keys to open. Even if they are the same key, two must be present to unlock the box. Kind of like a missile silo. Biggest problem: finding another user if it is 3 a.m.—and no audit trail.
- Place the password in an envelope in a clear box or behind a glass panel. The user must break the glass to access the password. It would be hard to hide the fact that it had been opened, but still no audit trail.
- Same as above, except add an alarm. Still no audit trail.

When I need to do this again, for my new systems, this is what I am going to do:

1. Create a second root account (a user with UID of 0).
2. Write a program that returns the second root account's password, but can only be executed by root. It also sends e-mail to the root user each time it is executed.
3. Set up the second root's account to allow access only from a fixed terminal. Use */etc/securetty* for this.
4. Use *sudo* to limit users who can execute the command, including the auditing features.
5. Write a *cron* task that executes *last* on the second root account every minute and sends e-mail to root when the user logs in.
6. Modify the second root's start-up files not just to load a shell but to

load a shell and then disable the account when the user logs out. Also, set up the environment for the account to record all actions the user takes.

This method allows one access for each secondary root account password and provides auditing capabilities and immediate notification. What it doesn't do is allow multiple accesses to the secondary account while the staff is unavailable. It also requires some very careful programming of the access program and start-up program. Some advice: Do not use a shell script, and use PGP if possible for the password being stored.

That's all for this time; please note the new e-mail address when sending me column comments and ideas. ◧

Christopher Curtin is the Team Lead for Server-side Client-Server development for Manhattan Associates. His e-mail address is ccurtin@mindspring.com.



CIRCLE 4 ON READER SERVICE CARD



CIRCLE 132 ON READER SERVICE CARD

The Quest for Ubiquitous DCE



by Jonathan Chinitz

After having worked with DCE for a number of years now, I am convinced of the following facts:

- DCE has no competition when it comes to a scaleable and interoperable security framework
- Building security frameworks is not for the light-hearted and should not be attempted by the average programmer
- Organizations could actually solve some serious security problems if they used DCE
- Companies want to "deploy" DCE but they do not necessarily want to know that it is "there"

Let me explain.

Who Wants to Buy "Plumbing"?

For DCE, the age-old motto "seeing is believing" just plain doesn't work. Because "seeing" DCE exposes some of its rough edges. It is, after all, an extensive framework of services and programming interfaces. "Seeing" DCE means programming directly with it. This requires knowledge, training, experience, and a software development strategy that says: we are building a new foundation for our business. Out with the old, or at least out with as much of the old as possible, and in with the new. How many companies do you know that are willing to undertake a major re-engineering, retraining and recoding effort without guaranteed success? And even with guaranteed success the cost might be prohibitive.

I recall a conversation I had a year ago with a colleague of mine employed by a large university, while we were in the lounge at one of the nation's airports. This individual has been a long-time proponent of the DCE technology. He told me that no matter what he did to convince the university administration that DCE was a real and viable technology for the university's business applications, he just couldn't sell it. (How on earth do you sell 1/4-inch piping when your audience wants Victoria's Secret middleware?) This person concluded that in order to sell DCE he had to show what it could do with commercial off the shelf (COTS) products that were of importance to the university. He needed to make DCE invisible to the naked eye. And that is exactly where good plumbing should always reside—underneath the applications, supporting the corporation's line-of-business applications, with very little fanfare, but with good solid results.

If You Build It and Bundle It–Will They Buy It?

One recent evolution in the DCE marketplace occurred when the DCE vendors were forced to rethink the whole notion of how DCE is priced and sold to the end user. For DCE to be a player in the Web marketplace, it has to conform to the business model of that marketplace: cheap, readily available, easily accessible, and bundled software. For those of us who have been involved in the DCE community for a while, these concepts were never voiced in the same sentence in a constructive manner.

Today DCE is bundled with IBM's AIX and OS/390 operating systems, with HP's HP-UX, and Digital's Digital UNIX operating systems. There is talk about

doing a port of DCE to Linux (I know because I raised the issue at the recent DCE SIG meeting in Bellevue, Washington) and making it free for everyone to use. Microsoft's Windows 95 and Windows NT ship with the DCE RPC as well. This has been true for quite some time, but people did not realize this until recently.

The reason is the press that is being generated around the ActiveX Core Technologies and DCE integration project that is being run out of the Open Group. This project will produce a reference implementation of the ActiveX core services (a fancy term for COM/DCOM) that uses both DCE RPC and DCE Security. When this work is complete (end of this year) you will be able to determine at run-time whether you want to run ActiveX components that use a DCE security framework—with complete interoperability across the enterprise. Don't be surprised if you see DCE showing up on your favorite desktop—you won't even know it is there.

So what's my point, you ask? I contend that even if you build the infrastructure, people are not going to buy it. Why? Because nobody wants to pay for infrastructure. The marketplace has decided that if technologies like DCE are to succeed, they must be free or bundled with the cost of the standard operating environment. Think about it: when was the last time you paid for your TCP/IP stack? But infrastructure alone is not a solution. Customers want to buy solutions. Customers need applications that make use of the infrastructure, all the time, without much ado about it. So what better way to catapult DCE into the "boxing ring" of large-scale applications than to tackle the fastest growing of them all: the Web.

DCE and the Web: A Match Made in Heaven

The DCE/Web Advanced Technology Offering from the Open Group (formerly the Open Software Foundation), promises to bring to the Web what is so direly missing from it, namely a scaleable security model with built-in authorization. With the DCE/Web, users can have protected access to their Web objects without having to dive into HTML or other languages to do so. They can accomplish all this with their COTS browser through a specialized proxy server that "tunnels" (read: encapsulates) HTTP traffic inside DCE RPCs. By doing so the proxy provides the data stream with the security properties that we have all come to expect from a standard DCE application. As a matter of fact, soon students at university campuses will be able to view their grades and register for classes safely and securely using this technology.

The Web was the perfect match for DCE-like technology. It lacked the basic requirements of any large-scale application, namely a robust security model, one that included not just encryption, but authentication and authorization as well; and a scaleable naming architecture, one that understands how to locate objects in a distributed environment and supports the federation of disparate namespaces through the use of junctioned mount points.

The Web also had the other magic component that made the integration of DCE possible: a proxy mechanism. This allowed vendors to develop the DCE infrastructure around the Web and "plug it in" through the proxy model. No changes to any desktop or server application code. Web technology also showed signs of immaturity when it came to scaling of the servers. This prompted Open Group to fill a finer niche: that of a scaleable, multi-threaded Web server.

It's Not Just the Web That's Out There

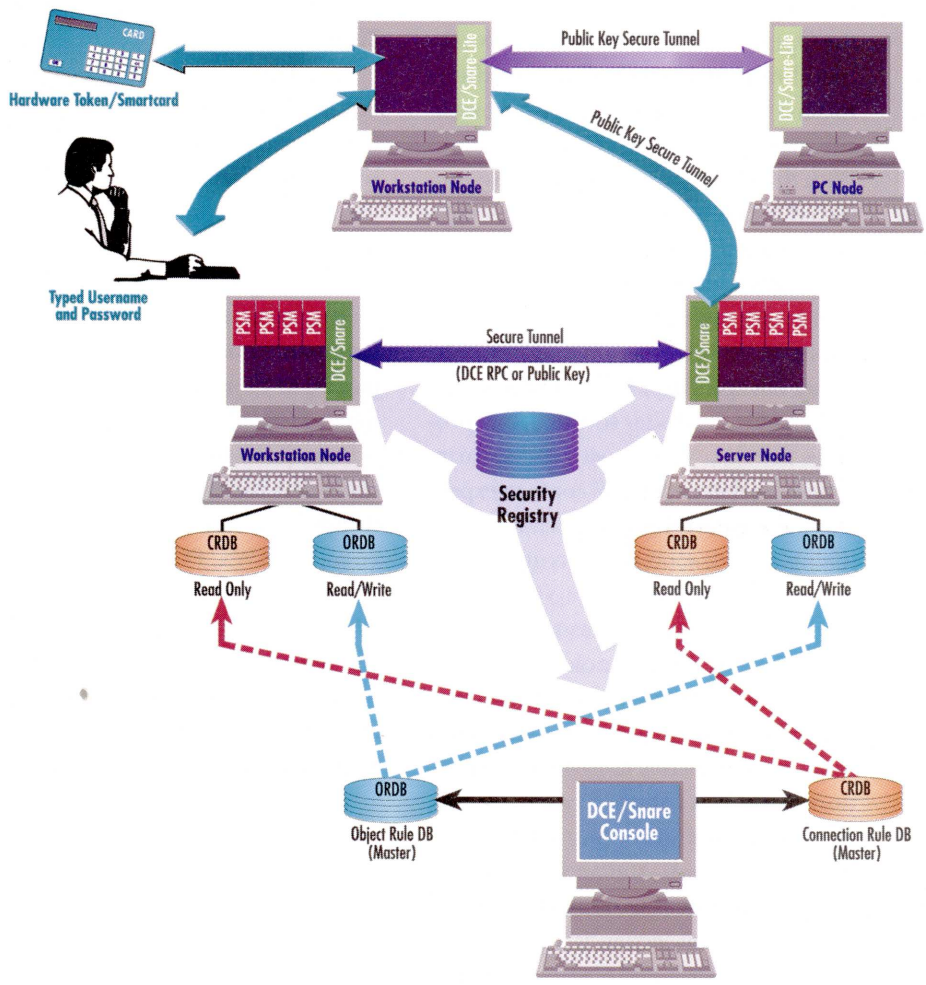
The Web is cool and the Web is hip, but the Web is not the entire computing industry. The Web, however, does highlight a broader market that DCE can tackle in much the same way: corporate Internets, sometimes called "intranets." Everywhere I turn these days I run into one of these intranet things. Companies have realized that for the past 10 to 15 years they have been investing in TCP/IP networks to support these off-the-shelf, ship-with-my-box applications that use these "Internet protocol things" such as TELNET, FTP, SMTP (mail), NEWS, POP, SNMP, and NFS. They now have a vast infrastructure of client-server applications, all communicating over the same network protocol stack, none of which share the same security model, much less have any security framework at all.

If only we could "DCE-ize" these applications then theoretically DCE would become ubiquitous. But doing so would mean replacing every COTS product with specialized DCE versions of it. This is not likely to happen. And what about the home-grown legacy applications that use TCP/IP through sockets? If we cannot bring applications such as TELNET and FTP to DCE, then why not bring DCE to (read: *under*) the applications? In the same way that the DCE/Web tunnels HTTP through a proxy, why not tunnel generic TCP and UDP traffic inside of a DCE RPC?

DCE/Snare

We at IntelliSoft have done just that. At the Decorum '97 conference this year

FIGURE 1 *DCE/Snare*



we demonstrated a technology that will bring the power of the DCE security architecture, and the promise of real single sign-on, to every TCP/IP-based application. We call it DCE/Snare (see *Figure 1*). Through a remote installation and configuration process, managed from a graphical console, you can drop DCE/Snare on any node in your network. The management console maintains a centralized database of connection rules and object rules. These databases are distributed to each and every DCE/Snare node automatically, thus enabling DCE to secure all TCP and UDP traffic to and from any machine. As an example, If you don't want to allow outbound telnet services from any or all machines, just tell the DCE/Snare and it will block them. The same goes for ftp transfers, Web accesses, news reading, etc. DCE/Snare can protect any application interaction that uses well-known ports.

On servers where sensitive corporate information is stored, DCE/Snare utilizes the DCE access control model to provide object-level protection and auditing of each operation and each access. This means that you can implement a uniform access methodology across your NFS mounted file systems and your FTP directories. Applications such as OpenView and Netview can securely manage your network nodes with SNMP, while a standard DCE ACL governs access to individual OIDs in a MIB.

Secure Single Sign-on

For the first time organizations will be able to implement secure single sign-on, with a universal network identity for each user. All that is required is that the user perform a DCE login and obtain DCE credentials. DCE/Snare will make these credentials available instantly to dozens of client-server applications across the corporate intranet. Since the DCE identity is universal, any data tunneled through DCE that carries this identity can be secured and authorized on any machine throughout the enterprise. This is all done with no additional programming, no library or DLL replacements, and no substitution of whole binary applications and databases.

The key to Snare is that it is transparent—as well it should be. You use

Article Reprints from hp·ux/usr

High-quality article reprints can help your company in many ways:

- Increased **EXPOSURE** for your product or service
- Credible, believable information that consumers **TRUST**
- Great **SALES** tools for trade shows, mailings and media kits
- Powerful **EDUCATIONAL RESOURCE** for consumers and employees



Contact Ruth Potter at

RMS
REPRINT MANAGEMENT SERVICES

147 WEST AIRPORT ROAD • PO BOX 5363
LANCASTER, PA 17606-5363
717.560.2001 • FAX: 717.560.2063
website: http://www.rmsreprints.com
email: sales@rmsreprints.com

the exact same applications that you used before, only securely this time. By using some of the standard features that are available in DCE today, Snare extends the sign-on to multiple authentication mechanisms that are not necessarily user name and password based (by adding a login registry attribute that is attached to your DCE principal). This allows for seamless integration of one-time password generators and even cryptographic smartcards into the sign-on framework, with no extra programming. Thanks to DCE's ability to support cross-cell secure RPCs, we do not have to limit ourselves to securing intranets. Using DCE/Snare in cooperation with the corporate firewall, system administrators can build a virtual private network of nodes, using the Internet as the connectivity medium. The nodes can be part of one security domain (read: cell) or several different ones. All access to the data is totally and transparently protected.

Snare of a Lighter Flavor

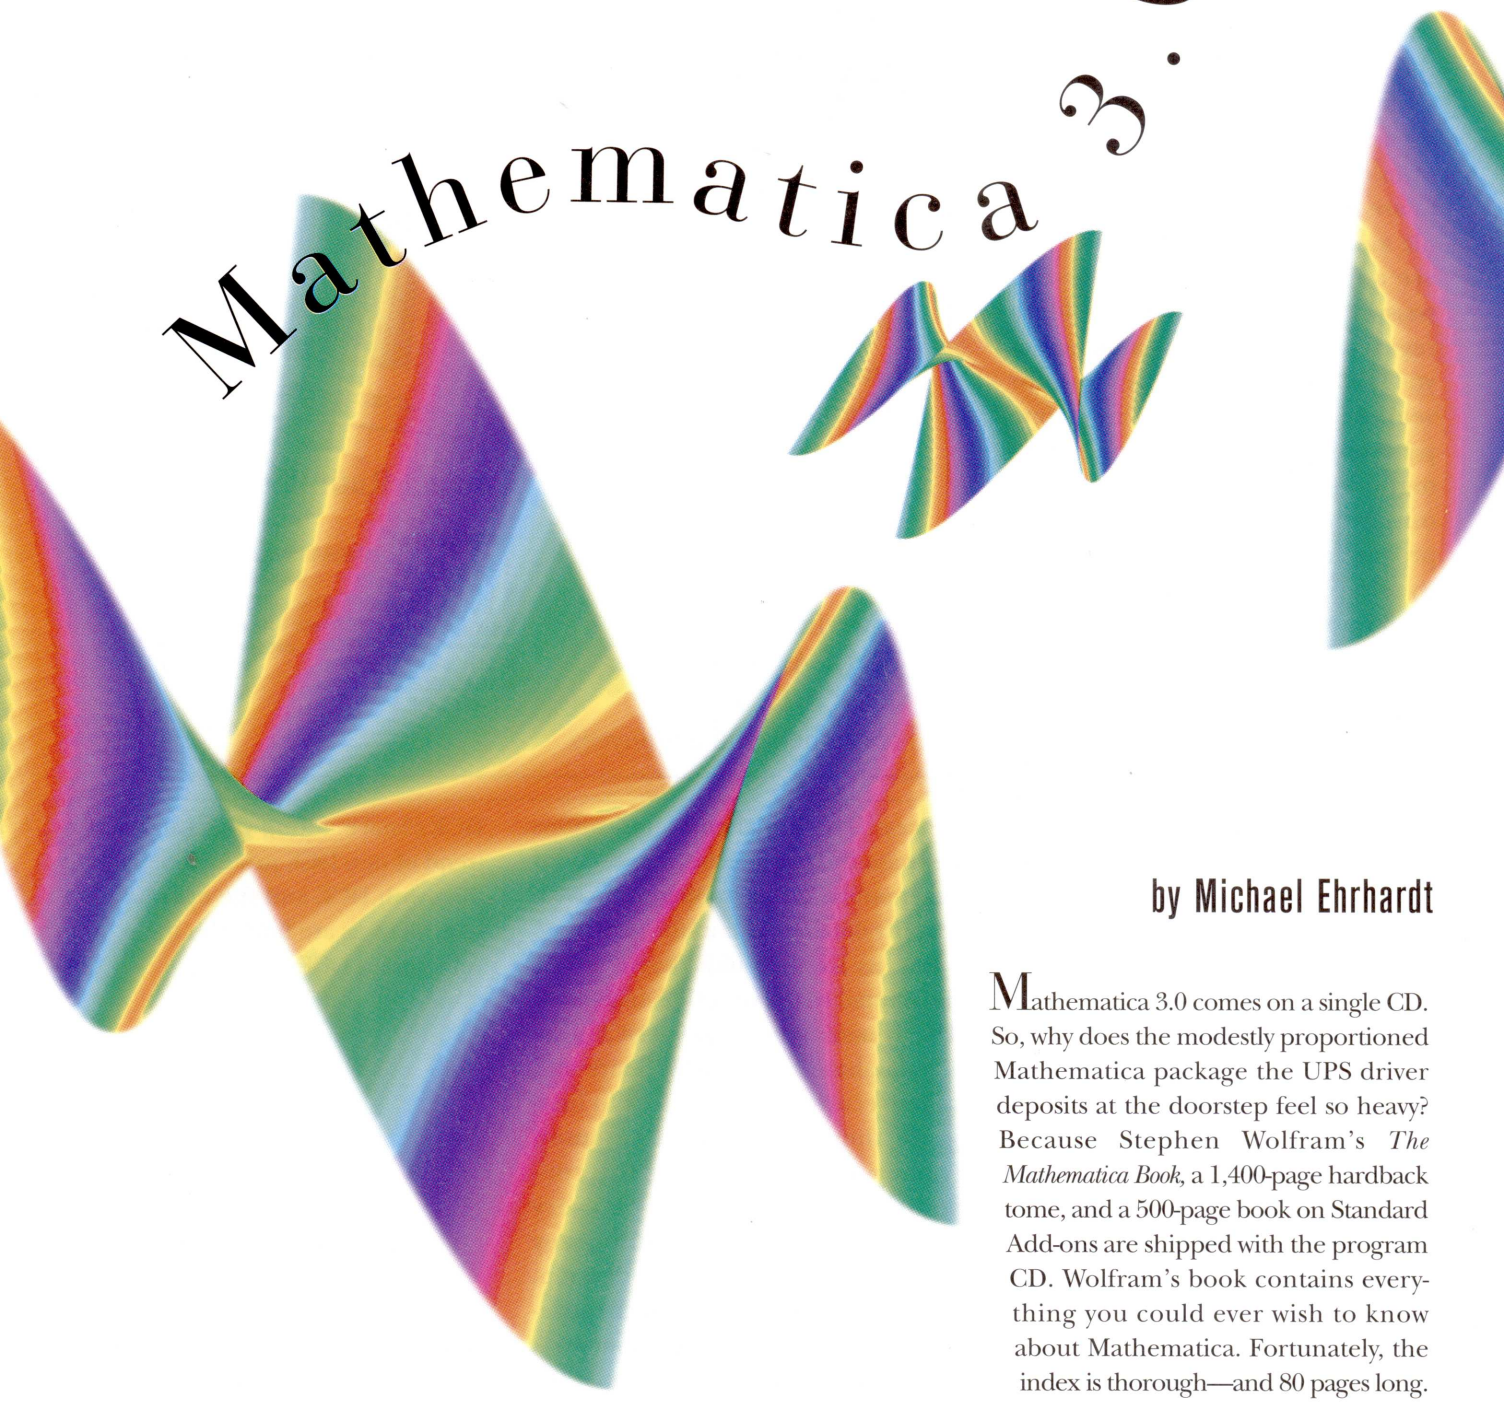
But what of those machines that do not have DCE capabilities? How do they interact with machines that are completely secured by DCE? Ideally we would want to provide users of these machines with a DCE security context that would be used whenever they attempted to contact a secure node inside the "Snared" environment. This would put them on equal ground with their DCE counterparts. This is the role of DCE/Snare-Lite. Snare-Lite enables a non-DCE machine to request DCE credentials from a DCE/Snare machine, using the DCE machine as its proxy login process. The initial communications between the two machines is secured using public key technology. The DCE security database acts as the certificate authority for the public keys that are distributed to the non-DCE machines. The dynamic schema of the DCE security database allows DCE/Snare and other security environments to integrate themselves into a common framework. By removing the constraint of requiring DCE on every desktop, but not compromising any aspect of the security of that desktop, DCE/Snare enables large corporations to roll out a security framework across different platforms, without the management overhead associated with DCE on every node.

Conclusion

It is my sincere hope that with technologies such as the DCE/Web and DCE/Snare, along with the current trend of new DCE products being bundled with DCE and the operating system, all of us are finally on the road to "ubiquitous DCE." ◧

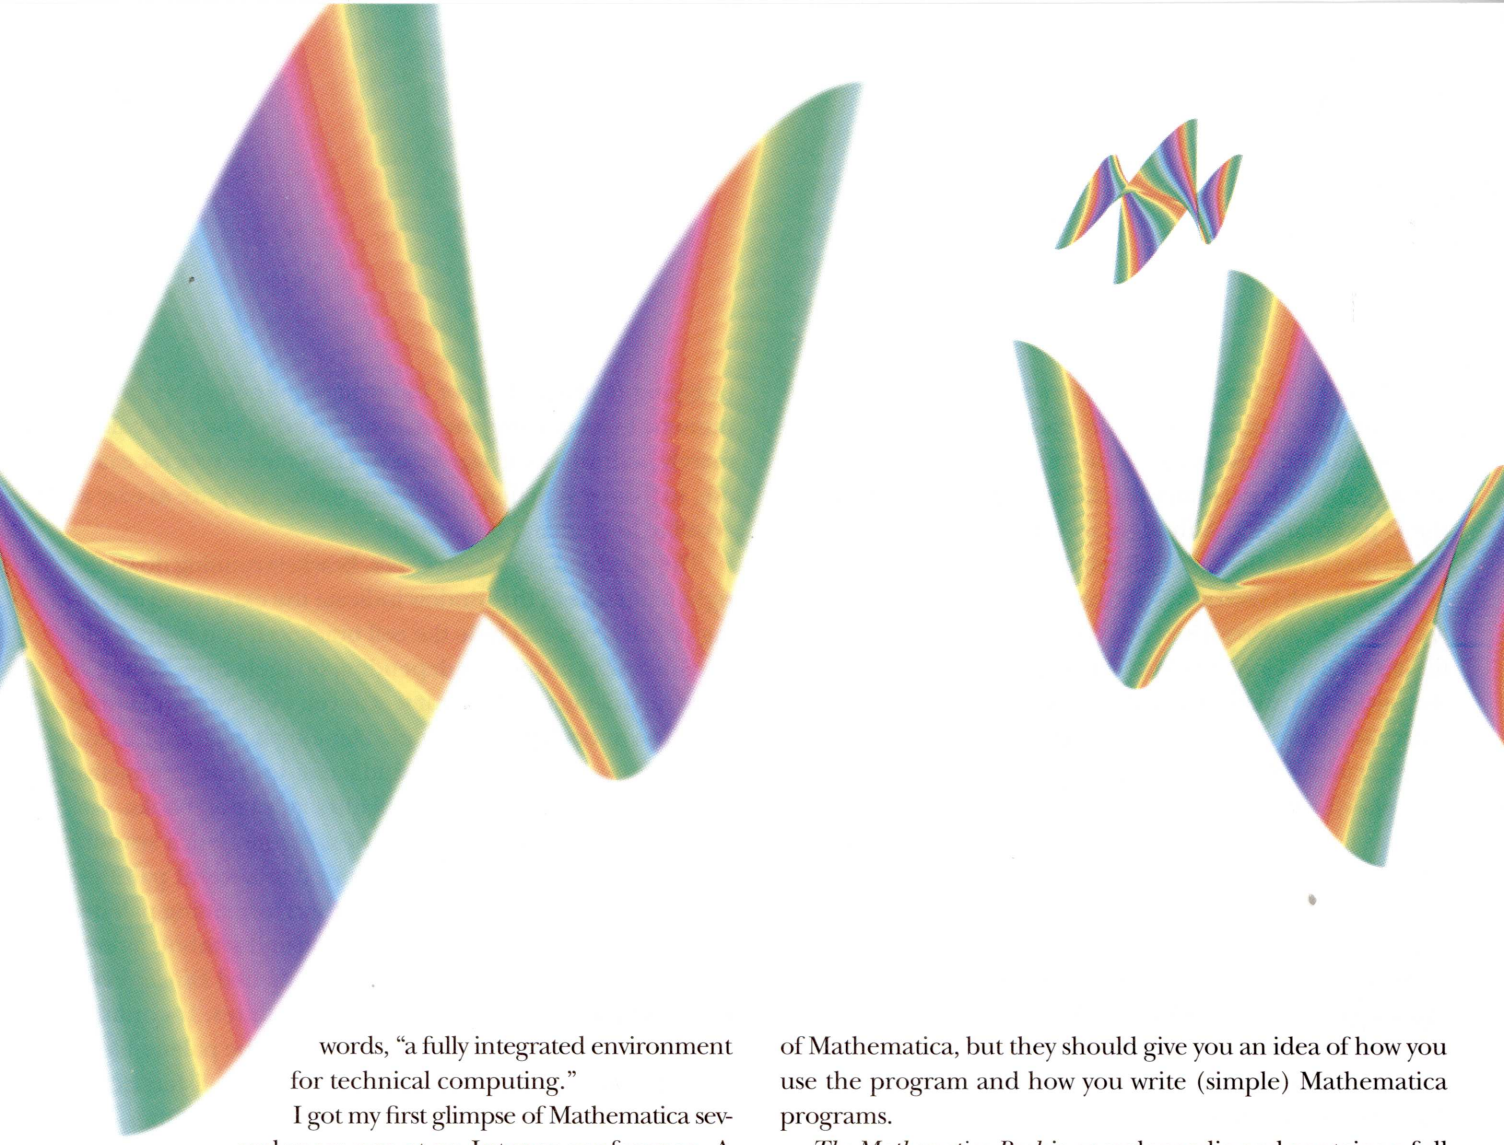
Jonathan Chinitz is president and founder of IntelliSoft Corporation, a maker of distributed management and network security products. He is a frequent speaker at industry events and has written extensively about DCE. In his spare time, he is chairman of the worldwide DCE SIG. He can be reached at jec@isoft.com. Information about DCE/Snare can be found at http://www.isoft.com/snare.

Mathematica 3.0



by Michael Ehrhardt

Mathematica 3.0 comes on a single CD. So, why does the modestly proportioned Mathematica package the UPS driver deposits at the doorstep feel so heavy? Because Stephen Wolfram's *The Mathematica Book*, a 1,400-page hardback tome, and a 500-page book on Standard Add-ons are shipped with the program CD. Wolfram's book contains everything you could ever wish to know about Mathematica. Fortunately, the index is thorough—and 80 pages long. As you look over the table of contents and browse through the chapters of *The Mathematica Book*, any notion you may have had that Mathematica 3.0 was simply a glorified calculator quickly fades. It is, in Wolfram's



words, "a fully integrated environment for technical computing."

I got my first glimpse of Mathematica several years ago at an Interex conference. A Wolfram representative was at the HP exhibition booth because Mathematica had recently been made available on the HP-UX platform. I was intrigued by the range and flexibility demonstrated. Also, the arbitrarily high numerical precision of Version 3.0 was a capability of even that earlier release. With the release of Version 3.0, I decided it was time to take a serious look at Mathematica.

This article is meant as an introduction for readers who know little or nothing about Mathematica but are interested in technical computing systems and may be considering buying the program. Mathematica can be used for any kind of calculating, both numeric and symbolic (and mixed); scientists of all stripes, engineers, science and mathematics educators all could find the program useful. Almost one hundred specialized commercial packages exist that provide specially designed Mathematica functions for different scientific, mathematical, and technical fields.

In a magazine article I can only scratch the surface of this large and very flexible program. I shall attempt to give an idea of what you can do with Mathematica by going through a series of relatively simple examples that I developed as I worked with it. These examples in no way demonstrate all the facets of Mathematica, but they should give you an idea of how you use the program and how you write (simple) Mathematica programs.

The Mathematica Book is encyclopaedic and contains a full reference section with a listing of built-in objects, with format and parameter examples. The Add-on volume describes in detail the additional functions shipped with the standard version of the program. But I found a couple of Mathematica books helpful in learning Mathematica programming, in large part because the wider range of more extended examples gives one a better feel for the basic programming techniques. Two books were especially useful: *Mathematica Graphics* by Tom Wickham-Jones and *An Introduction to Programming with Mathematica* by Richard Gaylord, Samuel Kamin, and Paul Wellin.

Wickham-Jones joined Wolfram Research in 1990 to work in R&D; he has been the main developer for graphics. His book is straight from the horse's mouth, as it were, and a very useful guide to graphics programming. Mathematica integrates graphics output with the flow of calculation—that's one of its most important features. You can manipulate graphics primitives in the Mathematica programming language. All the examples in this article involve graphics output.

The book by Gaylord et al is a good introduction to Mathematica programming techniques. The second edition includes exercises with solutions (some on the accompany-

ing disk). Their book demonstrates and encourages good Mathematica programming style. Basically, what you want to do is to take full advantage of the system's hundreds of built-in functions. This is a somewhat different style from that of, say, the C programmer. Although, if you know C, you will find much that is familiar in Mathematica syntax.

A third book I used is Alfred Gray's *Modern Differential Geometry of Curves and Surfaces*. The title says it all. What is unique about this text is that Gray uses Mathematica throughout to explore the intricacies, mathematical and graphical, of his subject. I must say, I think Mathematica is an excellent teaching and learning tool—always assuming, of course, that students have a firm pencil-and-paper grasp of the fundamentals.

Overview

In its standard form Mathematica 3.0 consists of two parts: The kernel, and the front end. The kernel evaluates expressions and does the actual calculating. The front end takes care of your input; it is your interface with Mathematica and handles formatting and graphics representation, and so on. You can run a command-line system without the front end, though I can't see doing it if you don't have to. On HP-UX, the front end is X Window System-based; the Mac and Windows 95/NT have their own front ends.

You can work in the front end without calling up the kernel. That is, you can edit text, change formatting, etc. in a document without recalculating anything. You can also run several Mathematica kernels, on local or remote systems, from a single notebook.

Notebooks

A Mathematica document is called a notebook. It is a platform-independent ASCII file that you can share with Mathematica users who have systems different from your own. Wolfram Research makes available a free Mathematica reader that enables colleagues who do not have the program to read your Mathematica notebooks. You can send notebook files as e-mail, on disk, etc., to other Mathematica users. You can also export notebooks as Web documents in HTML format, as TEX, and as RTF. Graphic elements can be exported as bitmap, EPS, metafile, and so on.

When you start Mathematica, the front end loads, and you are presented with a default new notebook window. All your input and output happens in the notebook. The kernel starts the first time you ask Mathematica to perform a calculation by typing [shift][enter], or by hitting the [enter] key on the numeric keypad. Using the main keyboard's [enter] key alone merely performs a carriage return. This enables you, for example, to format function definitions so that they are as easy as possible to read and understand.

The notebook is a series of cells. Brackets running down the right-hand side of the window indicate the extent and grouping of cells (you can hide the cell brackets if you wish). Cells have different properties depending on their type. Text cells allow you to enter and format text as you wish. You could write an article or book in a notebook, blending text, graphics, and calculations. Teachers should find it easy to create interactive course material in the form of Mathematica notebooks.

Figure 1 shows a screen shot of a Mathematica session. You see the notebook window as well as several palettes. Mathematica has a number of input and formatting palettes that make entering standard formulae and symbols as easy and fast as a mouse click. The palette labelled *Display* I created and added to the collection of palettes. It contains graphics options I use frequently. I placed the palette file in the Mathematica *palettes* subdirectory, so it is listed in the palettes menu along with the program's own palettes. It is easy to create your own palettes by using a menu option that converts a notebook to a palette.

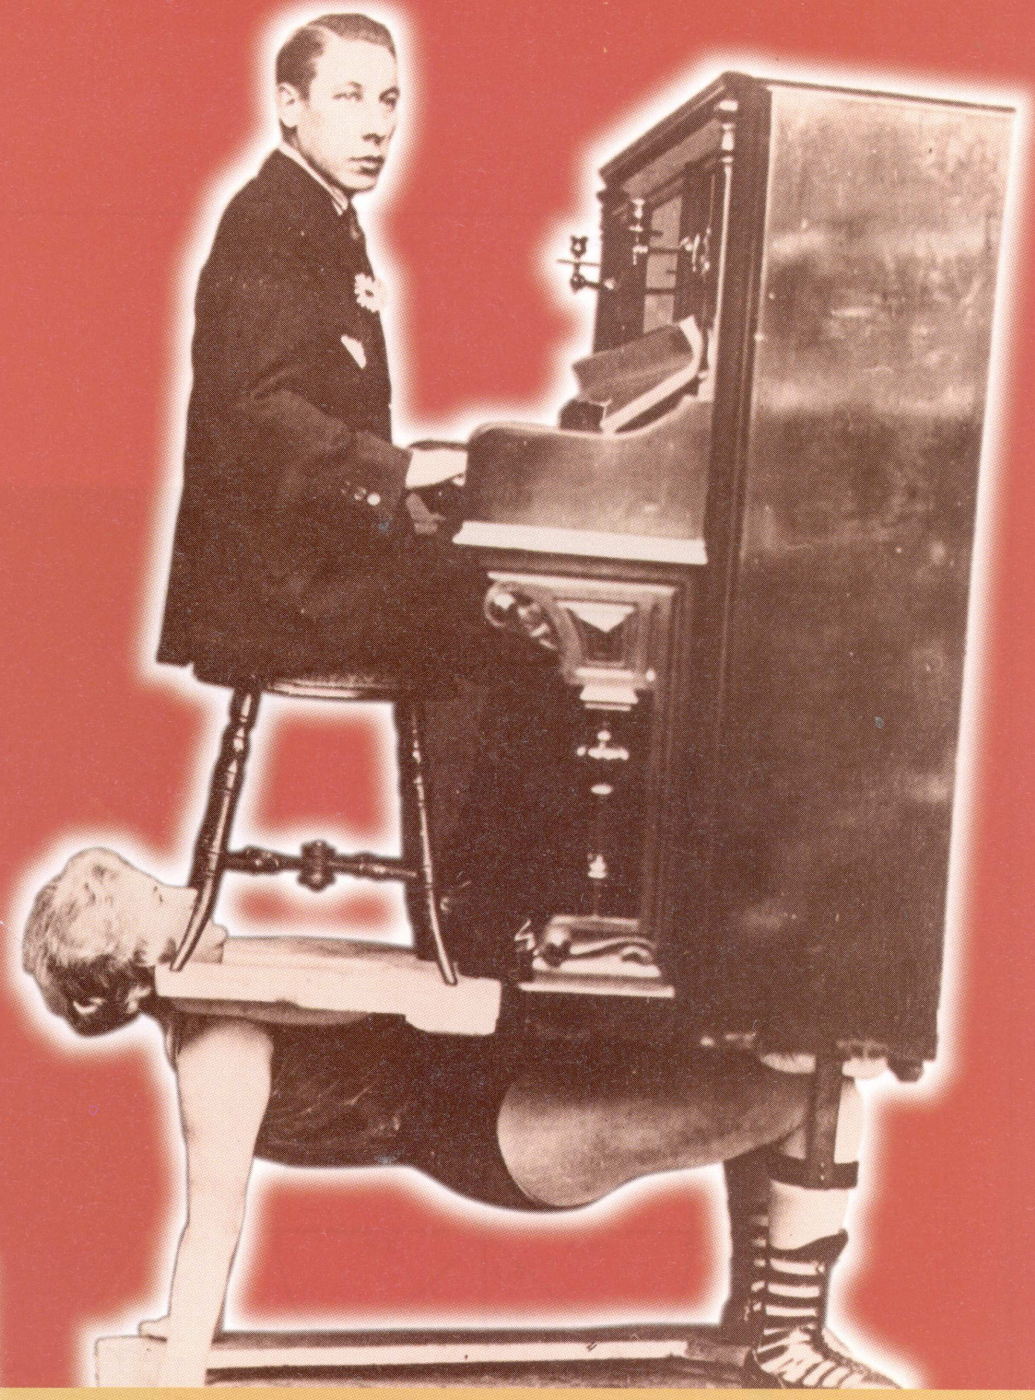
Basic formatting and cell manipulation options are available in the menus across the top of the Mathematica main window. The system has been designed to be easy to use. You do have to become used to the flow of cells, but going back and changing input and editing is easy. You can change parameter values in a calculation and recalculate simply by typing [shift][enter] again in the cell. Graphics objects are in PostScript format, and can be resized quickly with the mouse.

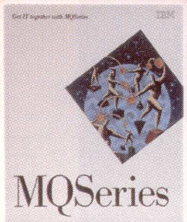
Entering input

Just start typing. If you type `3 + 4` and then [shift][enter], Mathematica responds by inserting `In[1]:` in front of your input and, on the next line, `Out[1]: 7`. You do not type the `In[...]:` and `Out[...]:`. Each session of a notebook starts with `In/Out[1]`, and increments by one until your last calculation. The % symbol stands for the last output, which you can use in the next input. You can also number the %, e.g., `%3`, to use output further back in the notebook as subsequent input. A menu option enables you to hide the `In/Out[...]` if you wish to.

The usual mathematical symbols have the usual meanings:

IBM





IBM software for HP-UX

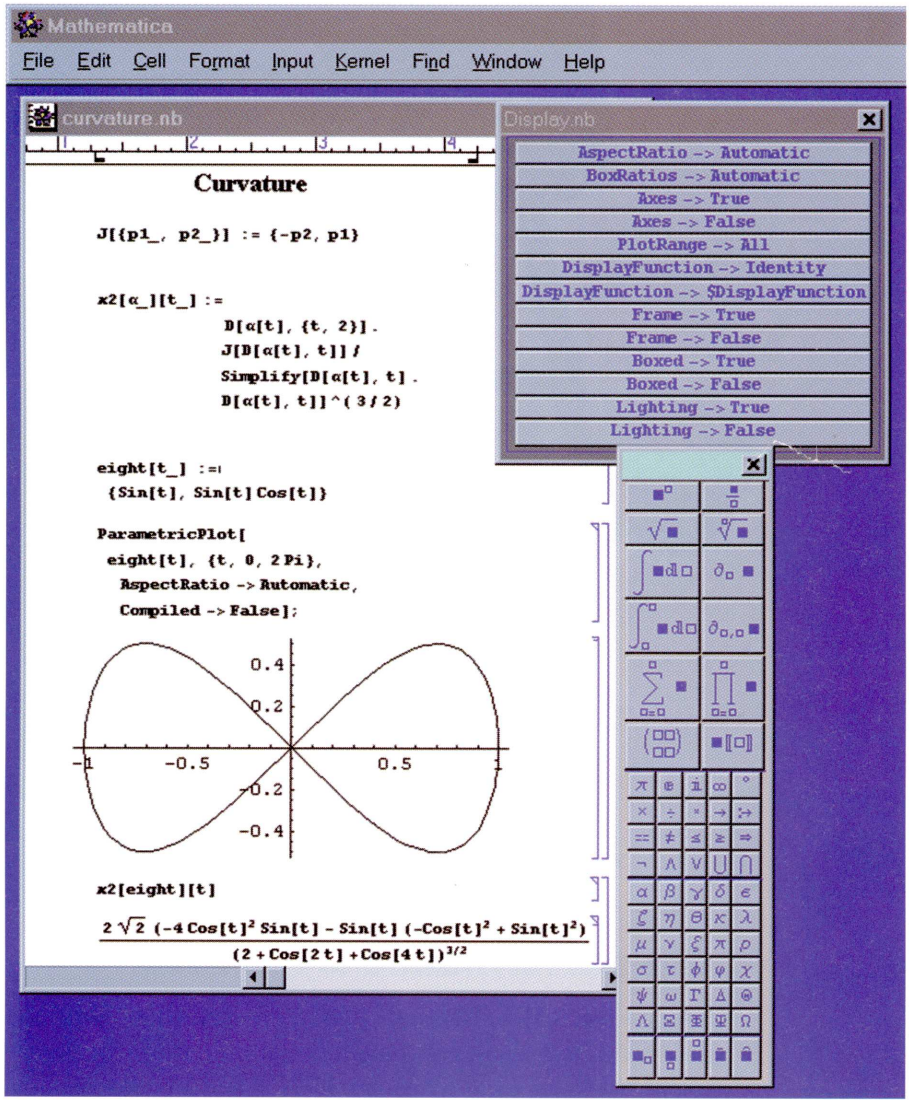
Take performance to new levels.

Strong-woman Mina Weezums (1883-1937) helped her accompanist hit impossibly high notes. Talk about providing full support – IBM's MQSeries for HP-UX can give you seamless integration between HP-UX and 20 other major computing systems. That means you can take HP-UX applications anywhere you want them to go. MQSeries is just one of many powerful and proven IBM Software solutions for HP-UX. Others include DB2, Net.Data, ADSM, Lotus Domino, Transaction Server and more. All of which will help performance go up and up and up. To find out more – including opportunities for resellers – contact your HP distributor or visit the IBM Software for HP-UX Website at **www.software.ibm.com/hp**

CIRCLE 158 ON READER SERVICE CARD

IBM, MQSeries, DB2, Net.Data, ADSM (ADSTAR Distributed Storage Manager), and Transaction Server are registered trademarks of International Business Machines Corporation. Lotus Domino is a registered trademark of Lotus Corporation. All other company and/or product names are trademarks of their respective companies. ©1997 IBM Corp. All rights reserved.

FIGURE 1 *A Mathematica Session*



- Plus +
- Minus -
- Times *
- Divide /
- Exponent ^

Operations have the same precedence as in standard mathematics. Multiplication can be, and normally is, indicated by a space between elements, rather than the asterisk. Thus `a*b` and `a b` both represent `a` times `b`. Moreover, a number followed immediately by a letter is taken to mean number times letter(s): `4ab` is 4 times `ab`, and you will notice that Mathematica spaces the `ab` the moment you type it right after the 4. Obviously, you cannot create variable names that begin with numbers (`x2` is all right).

Input and output can be formatted in several ways. The default, called StandardForm, is a blend of mathematical notation and computer language syntax that is consistent and unambiguous. Gaylord et al recommend using and getting used to the StandardForm style, and I think they are right. You can use TraditionalForm, which looks like math book text. This is useful for presentation, or, indeed, for writing papers and books in Mathematica. It is easy to mix the styles, using StandardForm for input, for example, and TraditionalForm for output.

You can use the mouse to place notation from a palette in your notebook; the cursor is automatically placed at the insertion point for the entry of input. So, for example, you can enter an integral in StandardForm:

```
In[2]:= Integrate[Log[1+x] / Sqrt[x],x]
```

or you can click on the integral symbol in the input palette (see *Figure 1*) and fill in the data.

To make Mathematica programming as intuitive and flexible as possible, the designers have made certain aspects of the syntax and form consistent. So, throughout Mathematica,

- Arguments to functions are enclosed within brackets—`func[a_, b_, c_]`
- Arguments within brackets on the left-hand side of function definitions have a following underscore, called a *blank* in Mathematica
- Names of Mathematica built-in functions begin with capitals—`Plot[...], Sin[x]`
- Except for common short forms such as `Sin, Cos`, and `Sqrt`, Mathematica function names are spelt out in full—`ParametricPlot3D[...], Integrate[...]`
- Lists are enclosed within braces—`{a,b,c}`
- Elements in a sequence are separated by commas—`f[a,b,c]`

You do have to get used to putting the underscores with variable definitions. Forgetting them is a common mistake. In addition to the single blank, two (`__`) stand for one or more arguments, and three (`___`) stand for zero or more arguments. You can easily write functions that take a variable number of arguments.

Parentheses are used only for grouping expressions—not to enclose function arguments, and not to enclose lists of objects. If you're used to writing C functions, just remember to use the brackets for functions.

The naming of functions in Mathematica is very intuitive—for the most part, everything is fully spelt out. That involves a bit more typing, but you really can't get the name wrong. There are more than a thousand built-in functions if you include the standard Add-ons. You don't have to remember them all: their names are what they do, run together as one word with each whole word part capitalized, as in `ParametricPlot3D, SurfacePlot, AddTo, NestList, StandardDeviation`, and so on.

Because Mathematica's built-in functions begin with capitals, it is best to name your own functions without capitalized first letters (at least) so that you can always tell at a glance whether a given function is one of your own or built-in. The capitalization obtains with all of Mathematica's basic constants as well:

- the natural logarithm base is `E`
- the square root of -1 is `I`
- π is `Pi`

It's a good idea to avoid single-capital-letter names as well; `N`, for instance, is the built-in command to convert expressions to real numbers and `D` is the derivative. If in doubt, you can always ask Mathematica by typing `?expression;` if `expression` is a Mathematica symbol, you will get back a description.

Functions that work in similar ways have the same form, so once you get a few basic function forms down, you can easily enter the arguments of a function you have not used before by

following the structure that is logical for that function. For instance, a basic function structure in Mathematica is

```
Function[expression, {iterators}].
```

So, we have `Table[expr, {iterators}]`, `Plot[function,{x, xmin, xmax}]`, `Integrate[function, {x, xmin, xmax}]`, and so on. The argument structure of these functions is the same. Once you get a feel for these basic structures, you can manipulate functions quite easily.

Expressions

The underlying and unifying idea behind Mathematica is that *everything is an expression.* The basic form for expressions is *head*$[arg_1, arg_2, \ldots arg_n]$. Arguments are numbered from 1 to *n* going from left to right and from -1 going from right to left, i.e., starting at the *nth* argument. The head is part 0 of the expression. You can extract the head of an expression either with the `Head` function or with the `Part` function. The head of an expression can be another expression.

The built-in function `FullForm` returns the standard expression form of its argument. This is a handy tool for checking the way Mathematica represents expressions you are using. For example, if we enter `FullForm[a+ 2 b]`, the output is `Plus[a, Times[2,b]]`. (Recall that the space between 2 and b means 2 times b.) That is how Mathematica represents the expression internally. The head of compound expressions is the name of the outermost function; in the above example `Head[%]` would return `Plus`. If we type instead `FullForm[a (b+2)]`, Mathematica returns `Times[a, Plus[2,b]]`. The head now is `Times`.

In its parsing mechanism Mathematica reduces expressions to combinations of the three basic types of expression called *atoms.* They are so named because, unlike other kinds of expressions, atoms have no parts. The three atoms are *symbol, string,* and *number.*

A symbol is a letter followed immediately by letters and numbers. Variable names, for example, are symbols.

A string is a series of characters enclosed within quotation marks—"This is a string."

The head of a symbol or string is the kind of atom it is. If we enter `Head["A test"]`, the output is `String` and `Head[variable]` returns `Symbol`.

A number is one of the four number types in Mathematica: *integer, real, rational,* and *complex.* The head of a number is the kind of number it is: `Head[6.009]` returns `Real` and `Head[7 - 3.1 I]` returns `Complex`.

When we write a Mathematica program, we create a set of rewrite rules that involve expressions. Mathematica uses pattern matching to evaluate expressions and apply transformation rules that, at the end of a series of substitutions for matched patterns, convert the component expressions to objects that can be calculated. The built-in functions are one kind of rewrite rule. The other is the user-defined rewrite rule.

Building Functions

A rewrite rule has two parts: a pattern on the left-hand side, and replacement text on the right. We can create a rewrite rule using either the `Set` or the `SetDelayed` function. The standard notation for `Set` is the equals sign (=). For SetDelayed we use colon equals sign (:=). The difference between the two is that `Set` evaluates the right-hand side and returns the value. We use `Set` for entering initial values and assigning values to parameters, etc.

`SetDelayed` does not evaluate the right-hand side immediately. We usually use `SetDelayed` for entering user-defined functions. The left-hand side of a function starts with a name. Names in the examples given below follow the naming/spelling conventions outlined earlier. The name may be followed by a sequence of patterns enclosed within brackets. The usual pattern is a symbol ending with an attached blank(s). Such a pattern is called a *pattern variable* or *labelled pattern.*

A pattern variable such as `x_` stands for *any* expression. If we make the function `cube[x_] := x^3`, we can enter any kind of number in `cube: cube[3/4], cube[Sin [Pi/3]], cube[1.2]`. In each case, the cube of the given number is returned. We can also specify the type of head a pattern must have to match input by appending a name to the blank. So, if we write `cube[x_Integer] := x^3`, only integer values will be evaluated. If we enter, for example, a rational, `cube[3/4]`, the output is simply the unevaluated function, `cube[3/4]`. It is worth pointing out that to Mathematica, rationals, integers, reals, and complex numbers are distinct, even though we can view rationals as a subset of the reals, reals as a subset of complex numbers, and so on. In writing transformation rules and conditional tests, it is sometimes necessary to specify the type of number involved.

The right-hand side of a function definition can contain a series of declarations, calculations, auxiliary functions, etc. These are separated by semicolons and enclosed in

Stuck using tar?

(or fbackup, cpio, or dump?)

What a sticky mess! Standard UNIX backup utilities force you to glue on scripts to make them work right, have performance like molasses, user interfaces that are clear as pitch, and reliability that could drag your career down into a pit. It's time to kick asphalt. Load BACKUP/9000 (it takes less than 10 minutes), and automate super fast, reliable backups and restores on any networked system via its slick user interface. Let BACKUP/9000 back up your Oracle databases hot, track tapes and files, manage media, schedule backups, etc. Don't get stuck - get something faster, easier, and more reliable. Contact us for a free demo today, before things get really messy.

ORBiT
GROUP INTERNATIONAL

+1-800-89ORBIT
+1-510-837-4143
Fax +1-510-837-5752
info@orbitsw.com
www.orbitsw.com



CIRCLE 72 ON READER SERVICE CARD

parentheses. Rewrite rules you enter in a session go into what is known as the *Global Rule Base,* which always contains the built-in functions. The declarations and names you use in compound functions all go in this base and remain there during the current session unless removed (by using the `Clear[...]` function). In general, we don't want dummy and temporary variables cluttering our Mathematica landscape, so it is a good idea to enclose the right-hand side of compound functions in the `Module` function. The syntax is

```
Module[{variables…}, expr; expr;. . .]
```

All the local variable names inside the function are listed at the beginning, within braces. A comma follows the list, so we don't need parentheses around everything. Simply write the expressions in the compound function, separating them with semicolons. A closing bracket ends the module. This way, the names you use inside the compound function are kept local to the function and do not appear in the Global Rule Base.

Transformation Rules

As we use our functions, we'll frequently wish to substitute a value for a variable, or set an option in a function. We use transformation rules of the form `a -> b (Rule)`, and `a :> b (RuleDelayed)`, often in conjunction with the replacement operators `/.` And `//.` The interesting thing about Mathematica's transformations rules is that they operate not only on literal expressions but on patterns as well. We can, moreover, give a sequence of rules—naturally, in a braces-enclosed list. Here are some simple examples:

```
In[1]: x + y2 /.  x -> 3
Out[1]: 3 + y2

In[2]: f[x] + f[y] /. f[t_] -> 2t
Out[2]: 2x + 2y

In[3]: f[x+y] - g[y] /. {f[s_]-> q, y-> r}
Out[3]: q - g[r]
```

The first is a straight substitution of `3` for `x`. In the second and third examples we have used a pattern: `f[t_]` means `f` of any expression, so `y` and `x` are replaced by `2y` and `2x`. Likewise, `f[s_]` means `f` of any expression, so `f[x+y]` is replaced by `q. y` goes to `r`. The single replacement operator `/. (ReplaceAll)` tries each rule in a sequence once only on each part of the expression. As soon as a matching pattern is found the result is returned and the part of the expression affected is no longer tested.

The multiple replacement operator `//. (ReplaceRepeated)` applies `/.` repeatedly until the result no longer changes. This makes transformations to subparts, for example, that would be missed by `/.` We must be careful with `//.` It is easy to create an infinite loop as in

```
x //. x -> 3x,
```

which will go until `MaxIterations` is reached (the default is 65536).

`Rule` evaluates the right-hand side immediately. That isn't always what we want. For example,

```
f[2,4,6] /. t_f -> Max[First[t],7]
```

yields an error: "*Nonatomic expression expected in position 1 in First[t].*" Why? Because `Rule` evaluates `First[t]` and `t` is not a sequence of elements. We want the rule to look for any expression with head `f (t_f)` and then substitute the Max of the first element in the expression and 7. The trick is to use `RuleDelayed`. If we write

```
f[2,4,6] /. t_f :> Max[First[t],7]
```

no error is generated. The evaluation of `First` is not made on `t` but on the sequence `2,4,6`, which is what we want.

With that condensed look at Mathematica fundamentals, it is time to look at a few real, albeit simple, programs. All the following examples involve graphics output, in part because of a personal predilection and in part because I wish to demonstrate some of the graphics capabilities of the program. I could have concentrated on statistics, on differential equations, on number theory, or on anything that lends itself to mathematical representation. The program is large and it is flexible.

Programming

We will begin with an old friend from high school geometry: the involute of a circle. If you unwind an imaginary string

around a circle, keeping it taut, that is, tangent at its point of contact with the circle, the end of the string will describe the involute of the circle. Given a circle of radius a centered at the origin, as the radius line (a units long) sweeps from some starting point through an angular distance ϕ, the length of string peeled off is $a\phi$ units. Our program draws the circle, the radius line to the point on the circle that corresponds to ϕ, the $a\phi$-long tangent line from that point to the involute, and the involute curve. We get not only the curve but the geometric elements from which it is generated.

The point on the circle corresponding to ϕ is (a cosϕ, a sinϕ). If you think of adding vectors, the point on the involute is the sum of the radius line at ϕ and the $a\phi$-long tangent line at ϕ. So, the point on the involute is at (a cosϕ + a ϕ sinϕ, a sinϕ - a ϕ sinϕ). We could simply use that in ParametricPlot to generate the involute, but instead we will create a list of graphics primitives and give them to the Graphics function:

Involute of a Circle

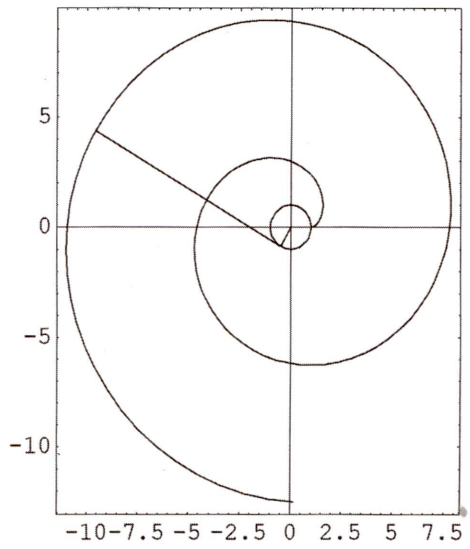
```
In[12]:= generator[a_, ϕ_] :=
          {Circle[{0, 0}, a],
           Line[
            {{0, 0}, {a Cos[ϕ], a Sin[ϕ]},
             {a Cos[ϕ] + a ϕ Sin[ϕ],
              a Sin[ϕ] - a ϕ Cos[ϕ]}}]
          }

In[13]:= Part[generator[1, ϕ], 2, 1, 3]

Out[13]= {Cos[ϕ] + ϕ Sin[ϕ], -ϕ Cos[ϕ] + Sin[ϕ]}

In[14]:= curve[a_, θ_] := Line[
           Table[
           Part[generator[a, ϕ], 2, 1, 3],
                        {ϕ, 0, θ, 0.1}]]

In[15]:= Show[
          Graphics[{generator[1, 10 Pi / 3],
            curve[1, 4 Pi]}],
           AspectRatio -> Automatic,
           Axes -> True,
           Frame -> True,
           PlotRange -> All ];
```



Our first function is called generator. It takes as arguments the radius of the circle and the parameter ϕ. The right-hand side of the function is a list of graphics primitives: it isn't really a function at all. Remember, lists are enclosed within braces. The primitives are Circle, and Line. These both have normal expression structure: name plus brackets with arguments. Arguments to Circle are the x,y values of the center (notice the enclosing braces—they're a list) and the radius. Line is very handy; it draws straight lines from one point to the next in the sequence of points you give as arguments. The whole (broken) line is a list of points, and therefore enclosed within braces, and the individual points are lists of x,y values, also enclosed within braces. generator gives us the circle and the two straight lines. What about the curve itself?

The points of the involute are the third argument in the Line primitive. We can extract them using the Part function. This also has a shorthand representation using double brackets, which we shall use later. How does it work? Well, the first part of generator is the Circle—we don't want that. We want the second part, Line. The outermost part of Line is a list, which we want. That is level 1 of the Line list. All of its parts are lists as well—the points of the line. Of those we want the third list, the points on the involute. So, going inside the levels of generator, we take the second element, Line (2), the first element of Line, the list of points (1), and the third point list (3). As you can see in the program, the syntax of Part is Part[expr, levels ...] all comma-separated. Hence, Part[generator[1, ϕ], 2,1,3].

Our second function, curve, creates the point list for the involute. It is a Line. The Line is constructed by building a table of point values. The Table command is one of the handiest in Mathematica. Table uses an expression and iteration parameters to create a list of elements. You can use it to build vectors, matrices, and tensors. The function evaluates the expression for each value of the iterator. You can have more than one iteration parameter, in which case the second parameter is run through for each value of the first. Notice the standard syntax of Table: an expression and a parameter list enclosed within braces. We use here the full iterator specification—parameter, starting value, ending value, and increment. If we had left out the increment, Table would have used 1. Line draws straight lines, so we want a small increment of ϕ so that the curve will be smooth. I could have given Table the actual formula for the involute, but I wanted to show the use of the Part command.

To draw the complete involute plus generating lines, we use Graphics, which creates a two-dimensional representation using the primitives and Show, which does the actual drawing on the screen. The structure of Graphics is Graphics[primitives, options]. Show can be used to redraw a Graphics object with different options or different values to options. It provides a quick way to alter and edit graphics output. Notice that the options are set using Rule: *option* -> *value* and that the primitives given to Graphics are created by evaluating generator and curve. I have given different values for ϕ in the two functions so that the involute goes to 4π and the generating tangent is shown at $\phi = 10\pi/3$.

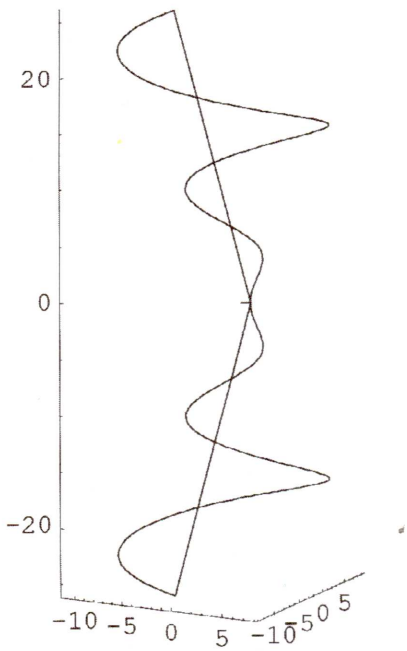
All the graphics functions take a fairly large number of options. Many we can leave at default values. AspectRatio we cannot because it isn't equal valued in *x* and *y* directions. The value Automatic makes *x* and *y* lengths equal. PlotRange is set to All to enforce a complete rendition of the curve. The default action is to reduce the picture to "the most interesting" area and clip out the rest. Axes and Frame do what their names imply. I have left the graph in its simplest form. We could add labels, tick marks, a frame name, and so on. We could change the font of the numbers if we wished. Now let's add a *z* value and draw the curve in three dimensions:

Involute Traced in 3D

```
In[16]:= curve3D[a_, θ_, z_] := Line[
           Table[ Flatten[{Part[
               generator[a, φ], 2, 1, 3], z}],
             {φ, 0, θ, 0.1}]];
```

```
In[17]:= spiral = curve3D[1, 4 Pi, 2 φ ];
```

```
In[18]:=
        Show[Graphics3D[{spiral,
          Line[{{0, 0, 0}, {1, 0, 0},
            Last[Part[spiral, 1]]}],
          spiral2 = spiral /.
                {q_, r_, s_} :>
            {q, r, -s},
            Line[{{1, 0, 0},
                  Last[
             Part[spiral2, 1]]}]
           }],
          BoxRatios -> Automatic,
          Axes -> True,
          AxesEdge ->
           {{-1, -1},
          Automatic, Automatic},
          PlotRange -> All,
          Boxed -> False,
          ViewPoint ->
            {1.174, -2.135, 0}];
```



Here we have created a second curve function, curve3D. The arguments to curve3D are the circle radius, the parameter, and z as a function of the parameter. Line is perfectly happy in three dimensions; we simply give three-valued lists as the points. Again, we use Table to generate a set of points. The new function here is Flatten. This function flattens out nested lists. The combination of the x,y values of the involute and the z values are within braces. But the x,y values extracted from generator are already a list within braces. Line wants a list of three values for each point; we can't have the extra braces (a nested list). Flatten removes the braces around the x,y values from generator so that x,y, and z values are one sequence enclosed in one set of braces. As you should be able to see now, Mathematica has a wide range of functions that manipulate expressions, allowing you to transform data in one context and use them in another.

spiral is the line generated by curve3D for the unit circle again, with the parameter going to 4π and z being 2ϕ. Notice that we are giving 2ϕ as an argument to curve3D rather than a numerical value. ϕ is evaluated as the table is constructed from generator, with ϕ running from 0 to 4π.

This time we give Show a Graphics3D object since our line is now a curve in three dimensions. Graphics3D is the three dimensional function that makes representations from primitives and options. We draw the spiral as well as a radius line from the center of the originating circle and a line from there to the end point of the curve. This last point we get by extracting the last point in the Line list in spiral using the Last function. Part is used again—remember the list of points is part 1. As you might expect, there is also a First function. Now look carefully at spiral2, suddenly popping up in the midst of the Graphics primitives. What is going on?

spiral2 is spiral with a transformation rule applied. Note the RuleDelayed symbol, :>. The pattern to be matched is any list of three elements—{q_, r_, s_}. Well, spiral is a list of three-element lists—the points of the curve. So the rule will be applied to each point in the list in spiral. The rule says take the negative of the third value, i.e., draw the curve with z values negative. A line is also drawn from the circle to the endpoint on the negative curve. The resulting graph is shown above. As you can see, transformation rules can make programming very concise.

The options are pretty obvious. BoxRatios is the 3D equivalent of AspectRatio. AxesEdge stipulates which edges of the bounding box axes are to be drawn on. The values must be 1, -1, or Automatic. 1 means an edge with a larger value of the coordinate, -1 an edge with a smaller value. Boxed set to False omits the bounding box around the graph.

The ViewPoint option sets the point in space from which the graphics object is viewed. It is relative to a three-dimensional box containing the object. The viewpoint is set using a special coordinate system scaled such that the longest side of the bounding box is length 1. The front end has a useful viewpoint viewer you can call up from the main menu. It shows a bounding box and a set of sliders with which to change the viewpoint. You can paste a value from the viewer into your notebook.

Having drawn a curve in three dimensions, let's turn our attention to surfaces. We shall start with Dini's surface:

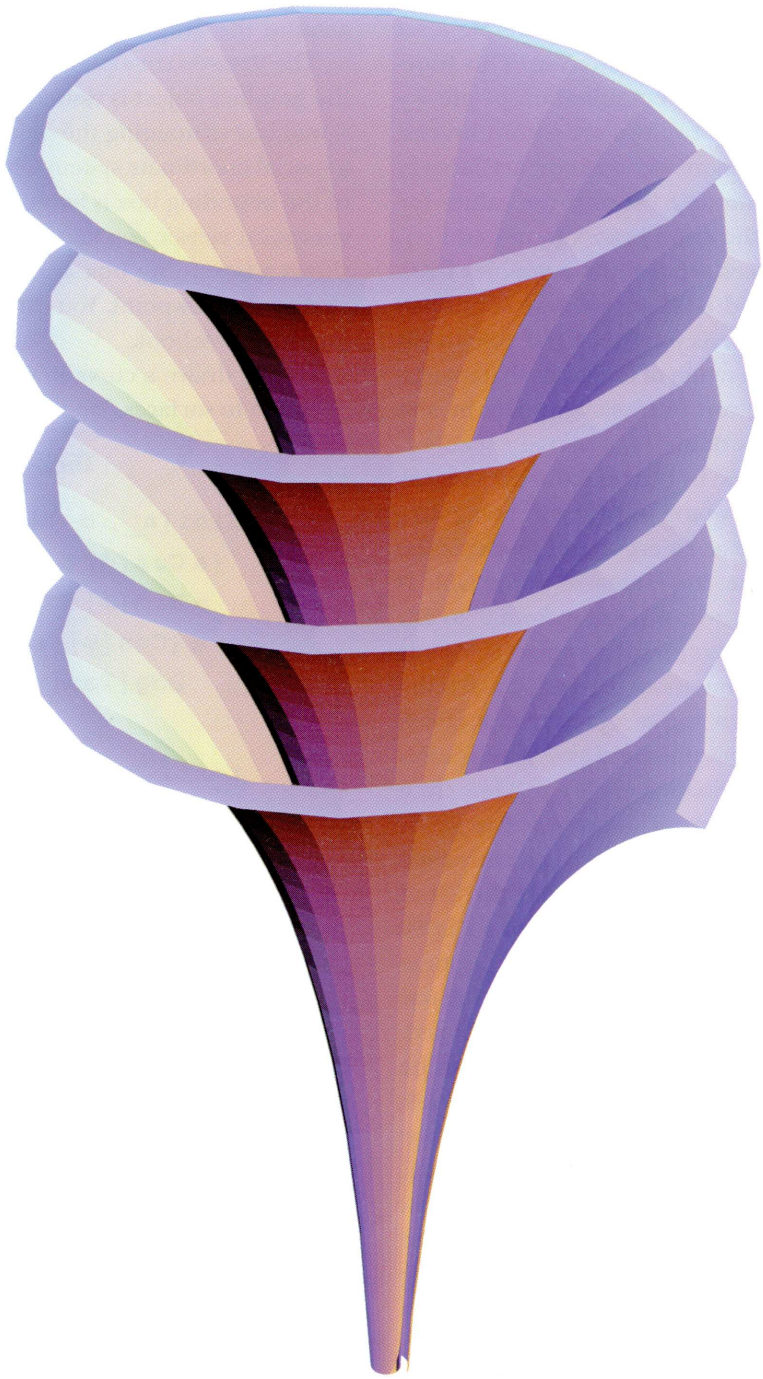
Dini's Surface

```
dini[a_, b_][ϕ_, φ_] :=
 {a Cos[ϕ] Sin[φ],
  a Sin[ϕ] Sin[φ],
  a (Cos[φ] +
    Log[Tan[φ / 2]]) + b ϕ}


test1 =
 ParametricPlot3D[Evaluate[
  dini[6, .5][q, r]],
  {q, 0, 8 Pi}, {r, 0.001, 2},
  PlotPoints -> {90, 50},
  DisplayFunction ->
   Identity];
Show[Graphics3D[{EdgeForm[],
    First[test1]},
    Boxed -> False,
    Axes -> False ]];
```

The graph of Dini's surface is shown in *Figure 2*. It is a twisted pseudosphere: a generalized helicoid of slant b constructed from a tractrix. Interestingly, it has constant negative Gaussian curvature of $-1/(a^2 + b^2)$. A tractrix is a plane curve starting at a point (a,0) such that the distance along

FIGURE 2 *Dini's Surface*



the tangent line from the curve to the vertical axis is constant for all points on the curve. A generalized helicoid is generated by a curve (here the tractrix) that is rotated about a fixed axis and translated in the direction of the axis with a constant velocity that is a multiple of the velocity of rotation.

Unlike our first example, Dini's surface is generated from a parametric representation. Thus, the right-hand side of the function `dini` is a three-element list of the *x,y,z* values in parametric form. The generating tractrix has the form `a{Sin[φ], Cos[φ] + Log[Tan[φ/2]]}`. The rotation parameter is ϕ. Note, again, that the list of values is enclosed within braces. The left-hand side of `dini` is interesting. Notice the double sets of brackets, one enclosing the constants, `[a_, b_]`, and one the parameters `[φ_, ϕ_]`. That's perfectly legal; in a sense it allows us to index the function by the constants.

`ParametricPlot3D` does just what it says. It creates a 3D surface plot using the parametric form—*{fx, fy, fz, color}*. *Color* is optional. It must evaluate to a real number between 0 and 1. I have not supplied a color function.

We feed `ParametricPlot3D dini`, with values for the constants supplied. `ParametricPlot3D` will plot a graph with ϕ running from `0` to 8π and φ from `.001` to `2`. We use the `Evaluate` command on `dini` because we want Mathematica to evaluate the function with the given constants and supply a point list that `ParametricPlot3D` can use with the iteration parameters. Unevaluated, `dini` doesn't produce a list of three expressions `ParametricPlot3D` can use to draw and the program will complain. Notice that the iterators are written in succession, enclosed within braces and comma separated, as always. As with `Table`, you give the iterator, the starting value, and the ending value.

The options follow the iterators. `PlotPoints` controls the number of sampling points. Here `{90, 50}` means 90 in the *x*-direction and 50 in the *y*-direction. `DisplayFunction->Identity` actually turns off the display of output from `ParametricPlot3D`. Why would we wish to do that? Because we are going to modify the output before displaying it. We give the polygons calculated by `ParametricPlot3D` to `Graphics3D`, using `Show`. Recall I said there was a `First` command. Here it extracts the polygon primitives—`Graphics3D` works with primitives. `EdgeForm[]` controls the look of polygon edges in 3D graphics. Called without parameters, it means no edges, i.e., a smooth unruled surface like the one in *Figure 2*.

Our next example is a shell drawn about a helix. It is similar to a tube about a curve, but in the case of the shell, the radius of the tube varies from point to point. You can generate some very interesting figures this way:

Helix Shell

```
Needs["ktbn`"]

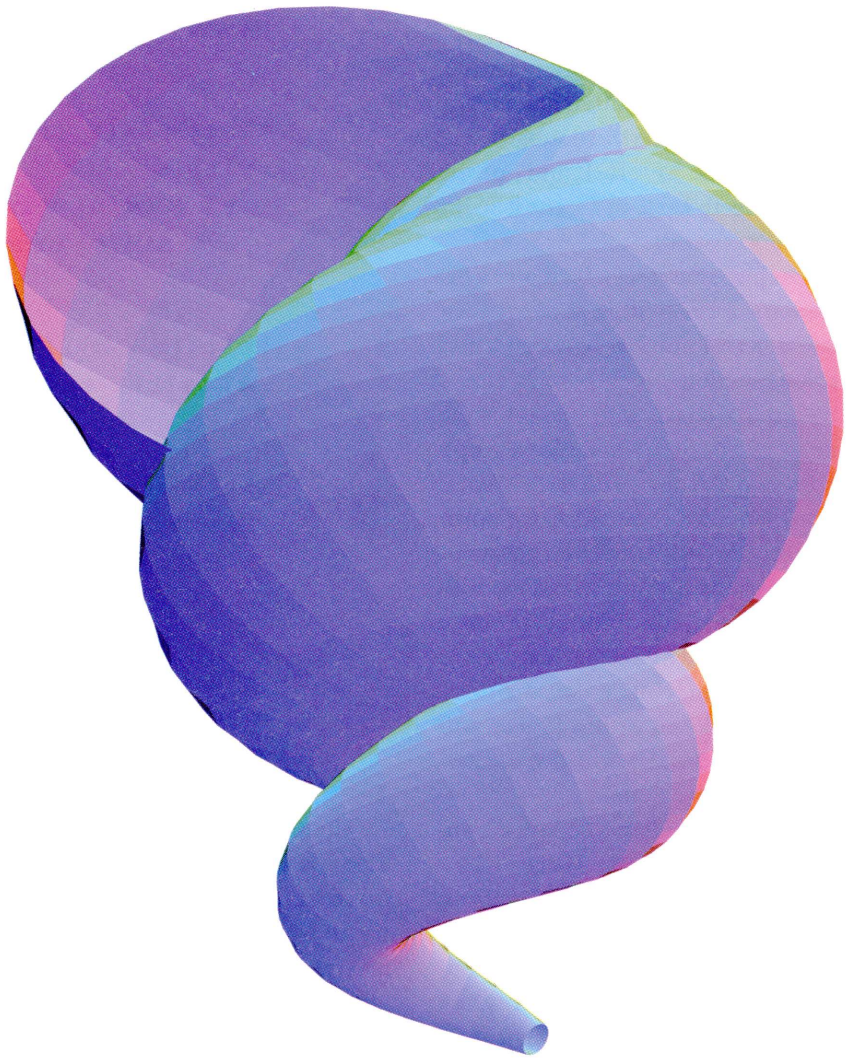

tubeshell[γ_][r_][t_, θ_] :=
 γ[t] +
  r t (Cos[θ] normal[γ][t] +
      Sin[θ] binormal[γ][t])


helix[a_, b_][t_] := {
 a Cos[t], a Sin[t], b t}


shell =
 ParametricPlot3D[tubeshell[
   helix[1, 0.5]][0.2]
     [t, θ] // Evaluate,
  {t, Pi / 4, 5 Pi}, {θ, 0, 2 Pi},
   PlotPoints -> {60, 40},
   DisplayFunction ->
               Identity];
```

The first thing you notice is the line `Needs["ktbn`"]`. The ` is a Mathematica context marker. I have created a file called `ktbn.m` that contains several lengthy differential geometry functions, in particular, **T**, **N**, and **B**, the tangent, normal, and binormal, which involve cross and dot products and derivatives. I used the lowercase names instead of the conventional capital letters so as not to clash with any Mathematica built-in objects. The front end has an option to save a notebook in what is called package form. Packages are just notebooks that group together related functions, constants, etc. Instead of having to write out the normal and binormal, etc., we simply read in the package `ktbn` using the `Needs` function. I placed the file in the *AddOns* subdirectory under Mathematica, where the program looks by default.

FIGURE 3 *Helix Shell*



tubeshell takes a parametric function for a space curve and adds to the x,y, and z values the value of the expression you see above involving the normal and binormal. Both normal and binormal are perpendicular to the curve γ.

All List objects have the property that f[{a,b,c,...}] returns {f[a],f[b],f[c],...} If you write {a, b, c} + x, Mathematica returns {a+x, b+x, c+x}.The helix function takes constants radius (a) and slope (b). We then give the helix function to tubeshell inside a ParametricPlot3D call, using Evaluate. Again, we use the option DisplayFunction->Identity, which turns off the display of the graph. We're going to modify it and the hand it to Show:

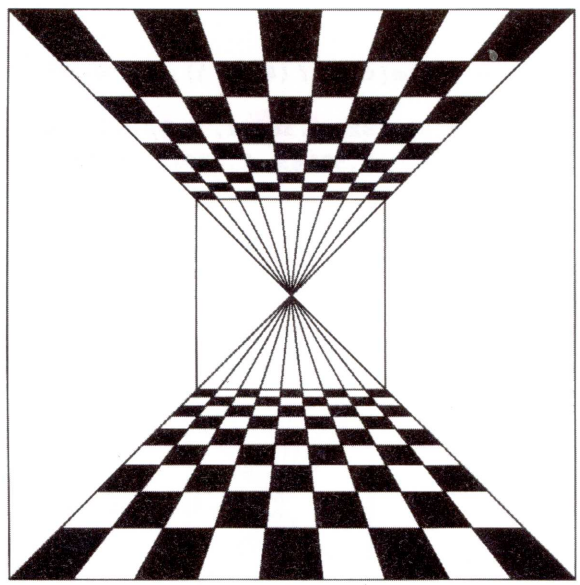
```
Show[
Graphics3D[{EdgeForm[],
  First[shell]}],
  BoxRatios -> {1, 1, 2},
  Boxed -> False,
  ViewPoint ->
  {-0.299, -2.431, -0.611},
  AmbientLight →
     GrayLevel[0],
  LightSources -> {
     {{0, 0, 1},
   RGBColor[0, 0, 1]},
      {{1, 0, 0.4},
   RGBColor[1, 0, 0]},
      {{0, 1, 0.4},
   RGBColor[0, 1, 0]}}];
```

We give Graphics3D the primitives generated in the ParametricPlot3D call by extracting them with First[shell]. As with the curve generated by dini, we use the EdgeForm[] command to suppress edge lines in the graph.

LightSources controls the direction and color of simulated light. This option can be set to enhance the sense of three dimensionality. The structure of the option is *{source1, source2,...}* where each source has a direction and a color. The direction is specified relative to a coordinate system with *x* and *y* directions horizontal and vertical in the plane of the display (or page) and *z* coming out toward the viewer in the positive direction. Only the *relative* magnitude of the components is relevant. The colors can be designated using GrayLevel, Hue, or RGBColor. I have used the last. The scale runs from 0 to 1.

The AmbientLight setting controls the simulated ambient illumination. Adding more ambient light lightens the effect of the colors. I have set it to zero using the GrayLevel function, that is, no ambient light, which is Mathematica's default setting. A sample graph is shown in *Figure 3.*

For our final example, we shall turn to a picture in Wickham-Jones's book:



He gives this as a demonstration of foreshortening and the convergence of parallel lines at the vanishing point. It is like an exercise in perspective drawing. He notes that the picture was produced in Mathematica but omits the code because it is "slightly complex." I thought it would be an interesting exercise to reproduce the picture in the book. My solution is a simple two-dimensional graph. I have added colors to the tiles, and generalized the picture in a function that allows one to alter the dimensions of the inner and outer boxes.

The graph above is square, which permits certain simplifications. Our function draws the outer and inner rectangles, the lines going from top and bottom of the outer rectangle to the vanishing point, and 9 × 9 grids of tiles on floor and

ceiling, alternating colored and white. What we do, basically, is use the two-point formula for a straight line repeatedly to get the vertices of the tiles. Let's stick with the square for the moment:

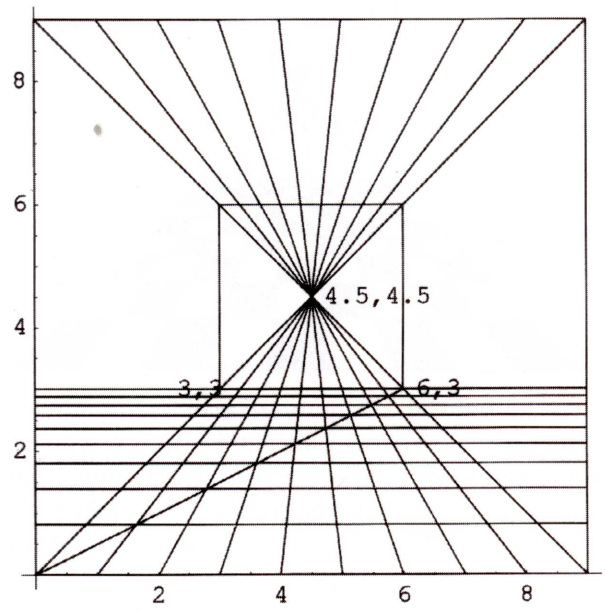
```
grid1 = Table[Line[{{i, 0}, {9 - i, 9}}],
                                {i, 0, 9}];

boxes = {Line[{{0, 0}, {0, 9},
               {9, 9}, {9, 0}, {0, 0}}],
         Line[{{3, 3}, {6, 3},
               {6, 6}, {3, 6}, {3, 3}}]};

directions = Line[{{0, 0}, {6, 3}}];


y = Table[4.5 j / (4.5 + j), {j, 0, 9}]

{0, 0.818182, 1.38462, 1.8, 2.11765,
 2.36842, 2.57143, 2.73913, 2.88, 3.}
```



The inner and outer boxes are drawn using `Line`. The outer square is 9 units by 9 units, the inner, centered within the outer at `(4.5,4.5)`, is 3 by 3 units. The vanishing point is in the center and is merely the point where the lines from top and bottom of the outer box cross. So, to draw the lines—which coincide with the edges of the tiles—we simply use `Table` to generate lines that go from bottom to top of the outer box: `(0,0)` to `(9,9)`, `(1,0)` to `(8,9)`, etc. That's `grid` in the code above.

Notice the line (`directions`) that goes from `(0,0)` to the right-hand bottom of the inner box at `(6,3)`. This is strictly a construction line: it goes through corners of tiles, giving us a way to calculate the *y* values of the horizontal edges of the tiles. We use the two-point formula for the straight line twice: for the line from `(0,0)` to `(6,3)` and for all the lines from `(j,0)` to `(4.5,4.5)`. The points where those lines cross give us the *y* values, so we simply substitute *x* values from one two-point equation into the other. In the case of this square, the equation is very compact, since the line from `(0,0)` to `(6,3)` is simply `x = 2y`. So our *y* values are: `y = Table[4.5j/(4.5+j), {j,0,9}]`. The generalized formula is slightly more complicated, but the same algorithm is used.

We apply the two-point formula again to create an array of *x* values. The *y* values work for all the tiles—it's a single list of numbers. For x values we need 10 sets of values—the 10 rows of x values for the vertices of the tiles. This will be a two dimensional array.

Mathematica under HP-UX

Two issues arise when dealing with Mathematica under HP-UX. The first problem is mounting the CD-ROM. The Mathematica CD-ROM contains distributions for several platforms, and the UNIX versions reside on the ISO9660 partition of the CD-ROM. One must mount this CD-ROM with the Rock Ridge Extensions. HP-UX's *mount* command is not capable of doing this, so the administrator must use *pfs_mount*, which ships standard with HP-UX 10.20, but can be downloaded from HP for users of HP-UX 10.10. Wolfram Research has a Web FAQ page that explains this issue in greater detail.

The second issue involves the use of fonts. Mathematica provides two different types of fonts for typesetting the special mathematical notation: Adobe Type 1 and Bitmap Distribution Format. It is our understanding that HP's implementation of X11 does not include the ability to scale the Type 1 fonts, so the BDF fonts must be used. The user must remove the Type 1 fonts from the X server's font path in order to prevent the occurrence of error messages during the startup of Mathematica. Also, if a user plans to display Mathematica 3.0 to a remote X server, the user must manually add the paths to the Mathematica fonts to the X server's font path.

—Technical Support,
Wolfram Research

The two-point formula is applied to the *y*-lines and the lines from `(j,0)` to `(4.5,4.5)`. Again, we use `Table` to create the numbers. We extract the *y* values from the list `y`, which we just created. Instead of `Part`, we will use the shorthand version, double brackets. That is, rather than writing `Part [y, k]`, we write `y[[k]]`. Our *x* values are

```
x = Table[j +y[[k]]-jy[[k]]/4.5, {k,1,10}, {j,0,9}]
```

We have all the information we need to draw the tiles using the `Polygon` primitive. The syntax is `Polygon[{{point}, {point},...}]`. The function joins the first and last points to close the polygon and *fills* with a color or gray level, which you can supply as we do in the final program. The tricky bit is figuring out how to set the iterators so that we draw every other tile, leaving white between colored tiles.

To make the writing a little shorter it is helpful to introduce the variable `n = Mod[k+1,2]`. We require two iterators, `k` for `y` and for the rows of `x`, and `j` for the values in the `x` rows. Our iterators will run as follows:

```
{k,1,9}, {j,1,8 + Mod[k,2],2}]
```

Notice that `j` runs in increments of 2: we're painting every other tile. To see how it works, let's look at the point for the lower left-hand corner of a polygon, the first point in the `Polygon` list of points:

```
{x[[k, j+n]], y[[k]]}
```

Remember that we are using the double bracket operator to extract parts of lists. When `k` is `1`, `n` is `0`, so for `j=1`, we get `x[[1,1]], y[[1]]`. `j` then goes to 3—remember that the increment is `2`. So the next point taken is `x[[1,3]], y[[1]]`. That's correct: we skip `x[[1,2]]`.

When `k` is `2`, `n` is `1`. So, the first point taken is `x[[2,2]], y[[2]]`. That's what we want: we skip the first tile in alternating rows so that we achieve a checkerboard tiling. And so on. The other three points of each polygon are constructed in the same manner. The ceiling tiles are the same as those on the floor but with *y* values starting at `9-y` and going down. The *x* values are the same.

The generalization to a pair of rectangular boxes is straightforward:

Perspective Function

pview generates a perspective grid with 9 rows by 9 columns of tiles. Arguments to *pview* are the height and width of the outer rectangle, and the width of the inner rectangle.

```
pview[l_, h_, box_] := Module[
  {mx, my, bx, by,
   x, y, view, boxes, n, inc},
  mx = l/2; my = h/2;
  bx = mx + box/2;
  by = my - (h/l) box/2;
  inc = l/9;
  boxes = {Line[{{0, 0}, {0, h},
          {l, h}, {l, 0}, {0, 0}}],
          Line[{{mx - box/2, by}, {bx, by},
          {bx, my + (h/l) box/2},
          {mx - box/2, my + (h/l) box/2},
          {mx - box/2, by}}],
          Table[Line[{{i, 0}, {l - i, h}}],
                      {i, 0, l, l/9}]};
  y = Table[
    my j inc/(j inc + (my bx/by) - mx),
                          {j, 0, 9}];
  x = Table[j inc + y[[k]] (mx - j inc)/my,
                  {k, 1, 10}, {j, 0, 9}];
  n = Mod[k + 1, 2];
  view = {Table[{RGBColor[j/10, k/12, k/10],
    Polygon[{
    {x[[k, j + n]], y[[k]]},
    {x[[k, 1 + j + n]], y[[k]]},
    {x[[k + 1, 1 + j + n]], y[[k + 1]]},
    {x[[k + 1, j + n]], y[[k + 1]]}}],
    RGBColor[k/10, k/12, (10 - j)/10],
    Polygon[{
    {x[[k, j + n]], h - y[[k]]},
    {x[[k, 1 + j + n]], h - y[[k]]},
    {x[[k + 1, 1 + j + n]], h - y[[k + 1]]},
    {x[[k + 1, j + n]], h - y[[k + 1]]}}]},
    {k, 1, 9}, {j, 1, 8 + Mod[k, 2], 2}]
    };
    Show[
    Graphics[{view, boxes}],
      AspectRatio -> Automatic,
      PlotRange -> All] ]
```

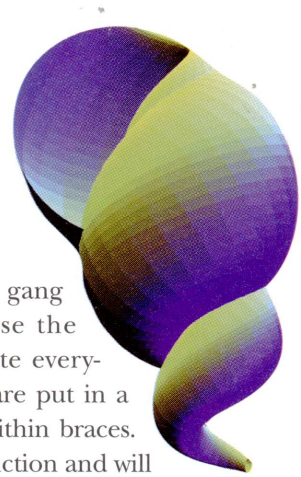
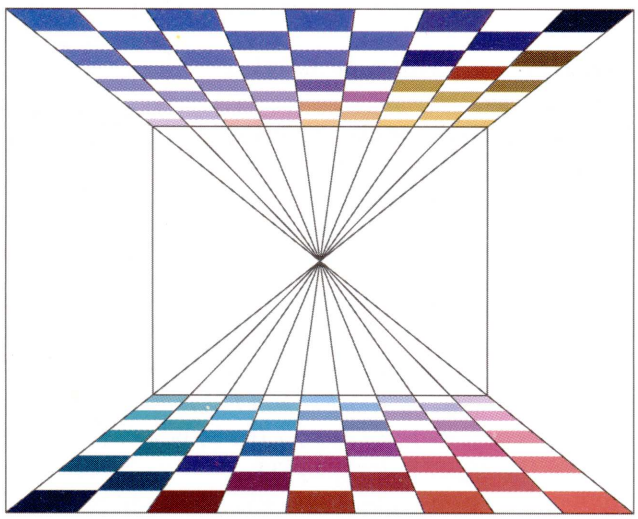
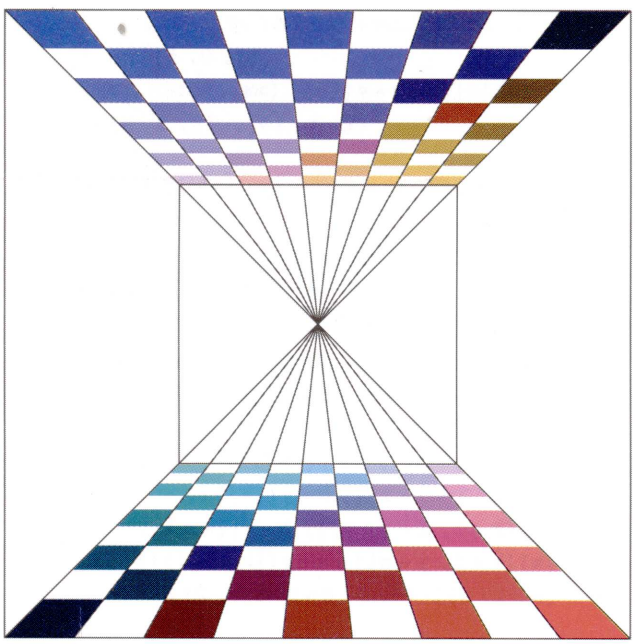


FIGURE 4 *Vanishing Point Boxes*





As you can see, we have a little gang of temporary variables. We use the `Module` function to encapsulate everything. The temporary variables are put in a list at the beginning, enclosed within braces. These will remain local to the function and will not go into the Global Rule Base. The function takes the width and height of the outer box and the width of the inner one as arguments. The inner box is centered and all the necessary lengths and point locations are calculated. Tables of *x* and *y* values are created using the two-point formula as we did for the square version. If we're feeling complacent, we can let Mathematica do the algebra for us. The `Solve` function will return the needed solutions. As an example, we can write

```
Solve[{y/(x-j inc) == my/(mx-j inc), y == k}, x,y]
```

and Mathematica returns the formula for the *x* values shown above, arranged only slightly differently. Notice the syntax for writing *equations*: the double equals sign is used (`==`). In `Solve` the equations are given in a list, and `x,y` means solve for *x*, eliminating *y*.

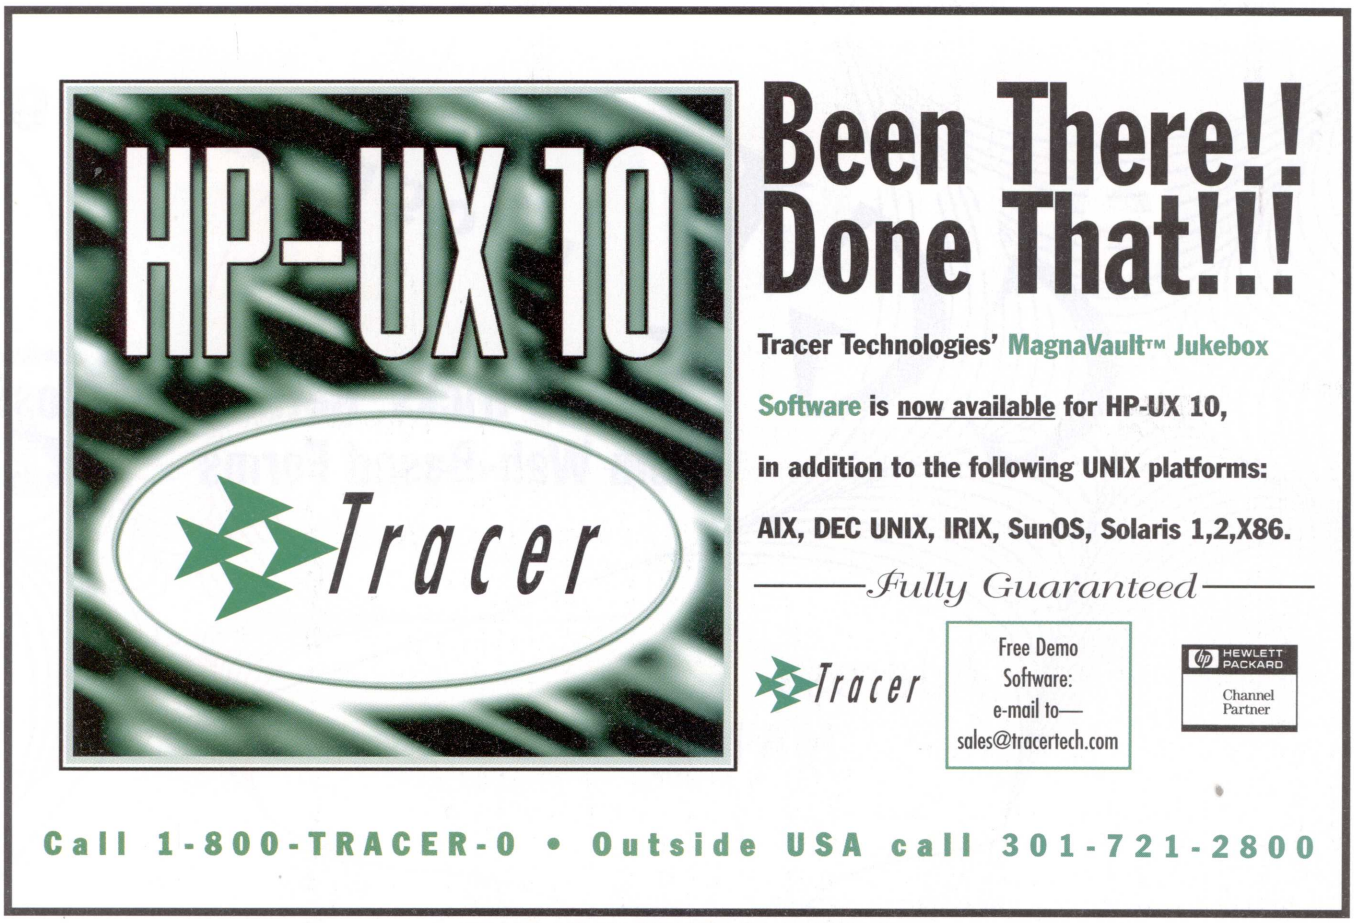
The polygons are drawn with different colors for each tile. That's the `RGBColor` function call. Sample results are shown in *Figure 4*.

Wrap-up

As I am a mere beginner in Mathematica, you can be sure there are easier, more elegant ways to do any of the foregoing examples. But I hope you now have a feel for how the program works and what you can do with it. I am sure that practice would bring proficiency and considerable speed—you can perform many manipulations in a compact fashion. The power of transformation rules and the ability to mix symbolic and numeric calculation greatly enhance the program's usefulness, and its flexibility encourages experimentation. It is an excellent educational tool for exploring the mathematical universe.

Like all number crunching programs, Mathematica will be perfectly content with the fastest machine on the block and loads of memory. The graphics capabilities of the program are extensive, but a 3D graphic with lighting and colors, etc., requires some serious calculating, and time. Well, that is why they're called *computers*.

Working in the Windows 95/NT version, I have experienced the occasional glitch when printing and sometimes in



CIRCLE 106 ON READER SERVICE CARD

trying to bring up the help browser. The main bug is a loss of window focus when the Show Page Breaks option is on. In fact, having that option on makes some files lock up the program and, sometimes, the system. That is a known bug, and developers will fix it at a maintenance release. For now, the best idea is to leave that option off. The bug does not affect the HP-UX version. For a couple of caveats on the HP-UX version, see the sidebar.

I enjoy working with Mathematica. It is so extensive that one could explore a wide range of mathematics topics and probably find oneself quite hooked on it. That is arguably better for the little gray cells than surfing the Net. The graphics output is especially pleasing. And don't think the program is limited to pure mathematics: the huge number of built-in functions and standard add-ons—not to mention a growing number of commercially available specialized packages—make Mathematica a useful tool in scientific and engineering computing. The interested reader will find Wolfram Research contact information in the At-a-Glance box. ◧

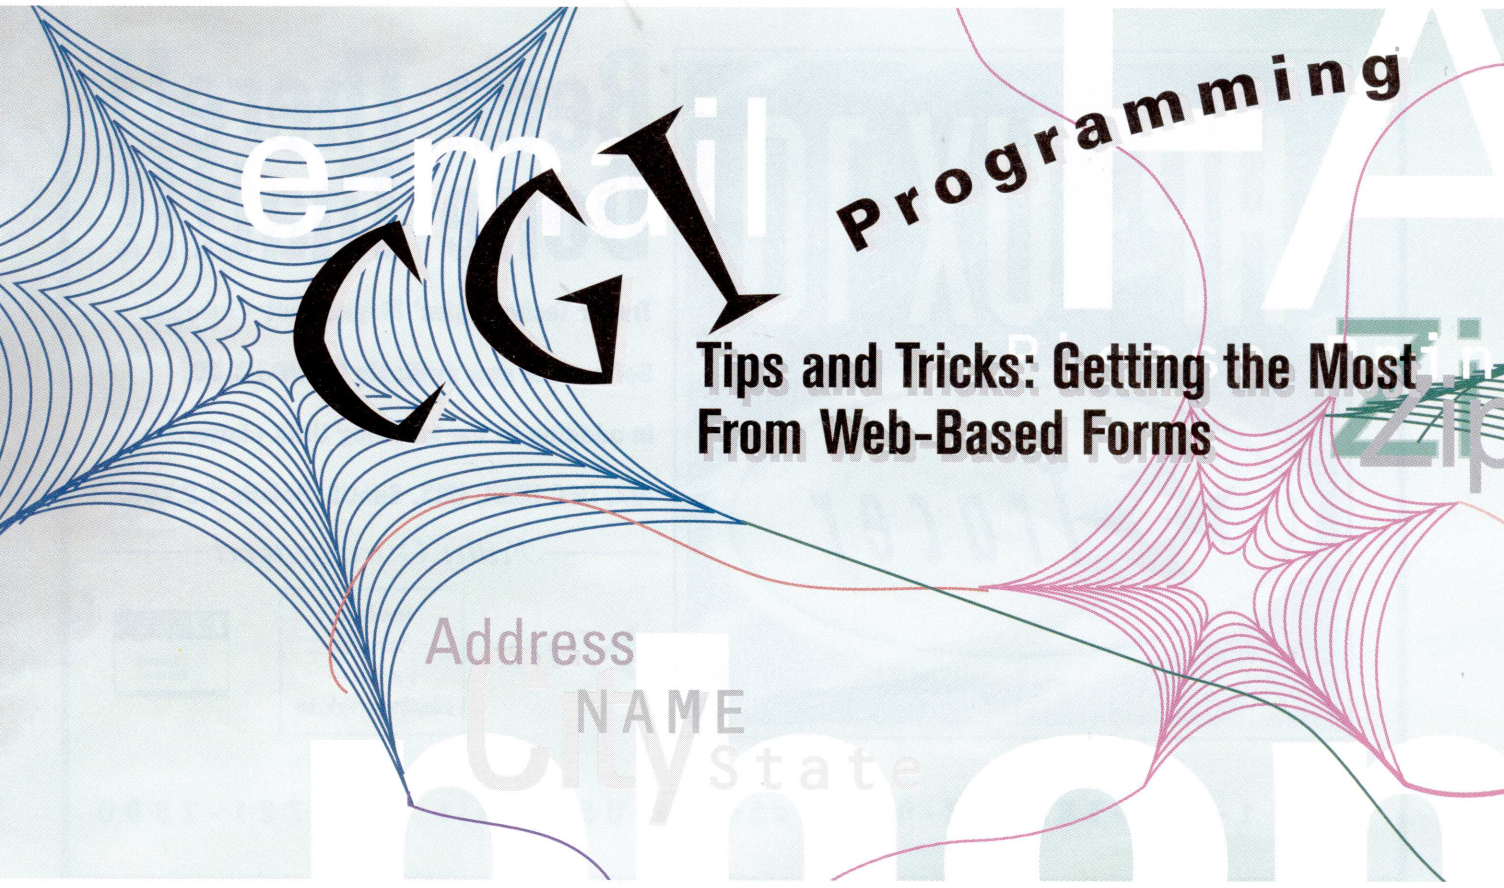
Michael Ehrhardt is managing editor of hp-ux/usr *magazine.*

At a Glance

Mathematica Version 3.0
Platforms:
HP-UX, version 10.0 and higher,
Windows NT,
Windows 95,
Macintosh,
Other UNIX platforms

Pricing:
HP-UX $2,495, $1,995 academic
Windows NT/95/ Mac/Linux $1,295, $895 academic

Technical Support:
Free tech support for first 90 days from date of purchase
Regular Service: $195 Windows/ Mac, $375 UNIX
Premier Service: $385 Windows/ Mac, $745 UNIX
Contact: *http://www.wolfram.com*
info@wolfram.com
1-800-441-MATH (6284)



CGI Programming

Tips and Tricks: Getting the Most From Web-Based Forms

by Chris Cobb

CGI is dead … **not**! As you read more and more hype each day about JavaBeans and ActiveX, don't forget that you can still get a surprising amount accomplished with "simple" CGI scripts. Okay, so this article isn't for everybody, but, if you're in the business of building Web applications, you may want to consider the most glaring omission in every CGI book I've seen to date: treating Web forms as single CGI scripts. Doing so eliminates the accompanying HTML files. Once you become used to this idea, I predict that you will be able to create much more complex forms without the accompanying design problems.

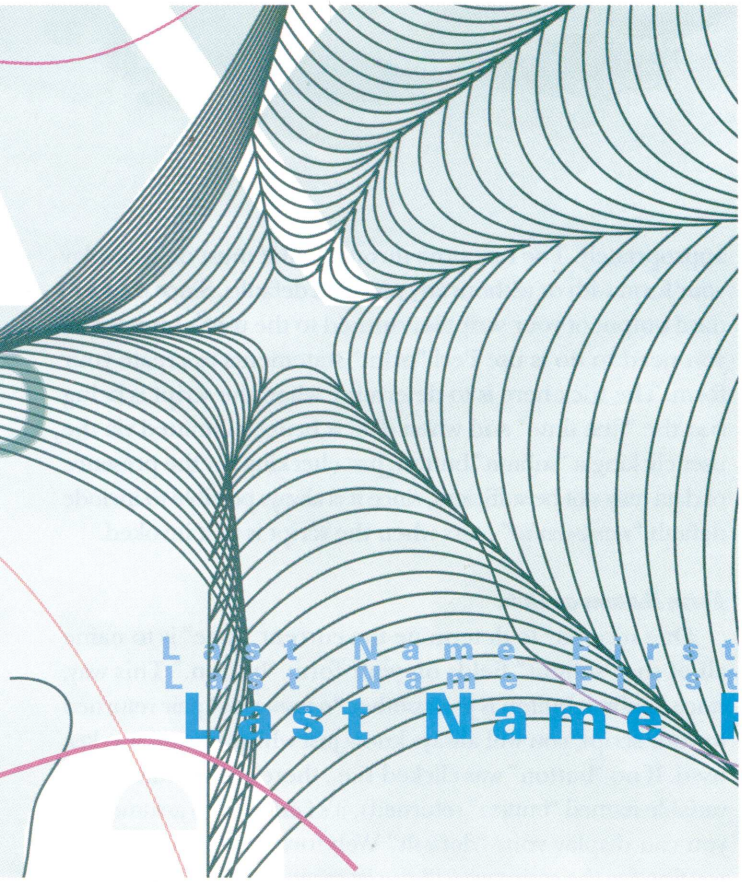
Assumptions

The examples in this article are presumptuous. They are about using Perl 5.002 or later along with a Perl library called *CGI.pm*. The scripts discussed here have been used successfully on HP-UX 9.0x and 10.0x using the NCSA, Apache, and Netscape Enterprise Web servers. While the *CGI.pm* tool discussed here will work on other platforms, such as NT and VMS, you won't find any mention of them (well, not after this paragraph anyway). However, though the examples are specific, the concepts presented are certainly compatible with any language, platform, and development environment you use. Any mention of browser specifics such as menus and buttons refers to Netscape 3.0.

This article also assumes that the reader is familiar with creating, installing, running, and debugging CGI scripts so no space is devoted to that. While it was tempting to include page after page of code listings and the resulting "screen snapshots," examples are minimal. Code appears only when necessary to help explain a point—the assumption is that the reader can access the Web, where full examples are available. If this assumption is unwarranted, feel free to contact me directly for copies via e-mail; otherwise, see the references at the end of this article for information on obtaining scripts that demonstrate the examples shown below.

Overview

This article has two parts. The main part covers the technical aspects of Web forms design; a secondary part discusses some graphic design elements. Just as desktop publishing combined many previously disparate and specialized fields into a single industry, the Web is combining DTP with programming, GUI design, database administration, etc. Web designers should have at least some understanding of all these areas.



The techniques outlined have been used with much success in several enterprise-wide Web-based applications developed within Hewlett-Packard over the last few years. Currently thousands of people within HP worldwide use these applications daily and, as of this writing, they are starting to see use by many thousands more outside of HP. The emphasis has been to develop some simple, effective tools and concepts that are easily reusable on many Web projects. While it may sound like yet more hype, experience has shown just how effective this really is.

Part I: Technical Design Issues

The CGI Environment

As you probably already know, CGI programming is kind of weird. Traditionally, computer programs start, open the necessary data files, interact with the user to modify data, close the files and stop. A Web-based form, however, merely displays a form and then stops. Once the user enters data and clicks a "submit" button, another process starts up to handle the data and then stops. The trick is to maintain some sort of "state" across multiple instances of your program in the "stateless" medium of the Web.

The Web server provides a little help maintaining state information across subsequent "browses" by providing Environment Variables that are available to your CGI scripts. Other ways to pass information into your applications include command-line parameters and HTTP "cookies," covered in more detail below.

Web servers bundle the information submitted through your Web forms and "encode" many of the characters before returning them to your script. This is a good reason to use a CGI "library" package that handles all of the character transformations for you.

Using Perl Libraries

There are several popular CGI libraries available for Perl programmers. The one discussed in this article is *CGI.pm*, created by Lincoln D. Stein at the Whitehead Institute Center for Genome Research in Cambridge, Massachusetts. It makes use of Perl 5's object-oriented capabilities. This well-designed library simplifies the life of a Web form designer and provides a terrific introduction to object-oriented programming in Perl. Another feature of *CGI.pm* is that it makes it very easy to migrate from the *cgi-lib.pl* library.

Creating Web Forms

You can greatly improve control over your CGI scripts by eliminating the *.htm* files associated with your forms and building them entirely through a single script. Once you dispense with using HTML forms that are separate from the CGI scripts that process the input, you will need to start designing your scripts with multiple modes of operation. While this may not sound compatible with "reducing the complexity" of your scripts, once you become used to this notion, you will soon have much greater flexibility in your Web applications. You could call this a rudimentary form of "event driven" programming because, as your forms get more elaborate, you will always need to anticipate which of many possible buttons were last used and, possibly, even where the user was prior to that. But let's not jump too far ahead!

GET vs. POST Forms

As with any programming, the first task is to figure out how to get data into your script. There are three basic choices:

- the command line
- environment variables
- reading files

(Oh, of course you are right—there are other choices such as shared memory segments, semaphores, sockets and signals, but these exceed the scope of this article!)

When your form is invoked through an "<a href>" URL, the

first time through will be an HTTP *GET* request. After that you can elect to use either the *GET* or *POST* methods. Why should you care? When using the *GET* method, form data is passed to your CGI script via the QUERY_STRING environment variable, whereas when using the *POST* method, form data is obtained by reading from STDIN. This means that using *POST* method forms allows you to have longer forms and pass much more data into your scripts than when you use environment variables. Depending on the operating system, Web server, and even the browser you use, the *GET* method can really limit your forms.

The *CGI.pm* library defaults to using the *POST* method. You can certainly use *GET* if you choose, but when using the *POST* method you won't have to keep checking to see if you have exceeded your system limits.

Retaining State Information

In the "stateless" medium of the Web, the next task is to decide how to maintain some sort of context across multiple Web "browses" with the CGI scripts you develop. Possibilities include:

- default form "action"
- form buttons as data
- self-referencing URLs
- command-line parameters
- fields and hidden fields
- environment variables
- HTTP Cookies
- data files

Which method you choose often depends as much on personal preference as on any technical reason. If, when reading some of the following examples, you find an affinity for a particular mechanism, use it!

As you read on, keep in mind that all of the following examples could easily be included in one single script. Remember that the script will perform different actions at different times as it runs with different data and in different modes.

Default Form Action

Without an accompanying HTML document containing the form data, it is necessary to run a CGI script directly, and you do this with an anchor tag just like any other URL (**). When the script is invoked, it needs a way to determine what "state" it is in so that it can respond appropriately. The first time through, all it must do is display your form with or without any prefilled defaults. Since the standard output of your script is returned to the user's browser, all you need to do is use Perl "print" statements to create your form. The trick here is to determine when the script is being run the "first time" and when data is being submitted via the user clicking a "submit" button. Just checking for the presence of data may not be sufficient, since it is always possible to include default "*name=value*" pairs when the script is first invoked.

Form Buttons as Data

One nice way to determine the current "state" is to name all of the "submit" fields on your form "button." This way, since a submit button is just another "*name=value*" pair returned to your script, you will always know just which button was last used. If no "button" was clicked (i.e., there was no value for a variable named "button" returned), it's fairly safe to assume that you can display your "default" Web form. Here is the main routine for the script we will use in examples throughout this article:

```
#!/usr/local/bin/perl -T
$Btn{'Submit'}  = "Submit Request";
$Btn{'Help'}    = "Help With Submittal";
$Btn{'ExitHelp'}= "Exit Help";

use CGI;
$editData = new CGI;
$Button   = $editData->param('button');
print $editData->header;  # Start HTML doc
if ($Button eq $Btn{'Submit'}) {
      if (&editFormData) {
            &endEmail;
            &displayConfirmation;
      } else {
            &displaySubmitForm;
      }
} elsif ($Button eq $Btn{'Help'}) {
      &displayHelp;
} else {
      # Default action:
      &displaySubmitForm;
}
print $editData->end_html; # End of HTML doc
exit(0);
```

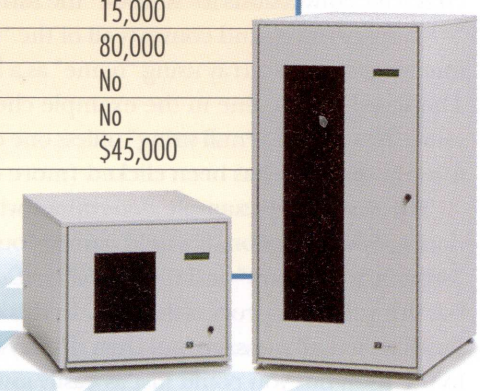
Thinking about DLT?

Think again! We know that many of you are thinking about DLT tape libraries. But recent advances in **8mm and 4mm** technology—and tape libraries made by StraightLine—might make you want to think again.

Model	StraightLine SL-400 Series	StraightLine SL-800 Series	ADIC Scalar	ATL ACL 4/52
Drive format	4mm	8mm*	DLT 4000	DLT 4000
Number of tapes in library	18 or 24	20, 50, or 150	18 to 52	52
Throughput	up to 25.8GB/hr	up to 108GB/hr	up to 43.2GB/hr	up to 43.2GB/hr
Total capacity	up to 586GB	200GB to 7.5TB	2.08TB	2.08TB
Bar code reader	Yes	Yes	Yes	Yes/6-digit
Number of drives	1 to 3	1 to 5	2 to 4	2 to 4
Head life (hours)	>20,000	>20,000**	10,000	10,000
Media uses (passes)	20,000	20,000	15,000	15,000
Drive MTBF (hours)	200,000	200,000	80,000	80,000
Tape drawer(s)	Yes	Yes	No	No
Removable tape boxes	Yes	Yes	No	No
List price starting at	**$7,995**	**$13,995**	$17,995	$45,000

*Exabyte's recommended 8mm tape—Exatape 170M Advanced Metal Evaporated tape
Sony's recommended 8mm tape—SDX-T3N 170M Advanced Metal Evaporated tape

**Exabyte Mammoth head life—20,000 hours
Sony AIT SDX-300 head life—30,000 hours



TAPE LIBRARIES

Led by Sony, HP, and Exabyte, tape drive technology has taken a giant step forward in reliability, performance, capacity, and price. Consider this:

Reliability — StraightLine has taken the best in drive technology and engineered tape libraries with mainframe-class robotics. Our unsurpassed accuracy—to .001 of an inch—means you'll get 'round-the-clock dependability. And that's not all.

Capacity and Throughput — The SL-800 series offers a range of **8mm** systems designed to meet your current needs and easily expand to meet your future needs, too. Choose Sony or Exabyte drives, then select the number of tapes—20, 50, or 150 tape systems—offering up to **7.5TB** of data storage and **108GB** per hour throughput.

The SL-400 series is now available with the latest in **4mm** technology—**DDS-3**. Using a StraightLine tape library, you'll realize **586GB** of total capacity and a data transfer rate of **25.8GB** per hour. Systems are available in 18 and 24 tape configurations.

Industry Standards — StraightLine libraries are compatible with all standard systems and software—Sun, HP-UX, SGI, IBM RS/6000, Sequent, NT, DEC; Novell, LAN Manager; Arcada, Legato, ArcServe, and many others.

Better than DLT! — If you're thinking about a tape library, think about this—we offer the best in reliability and performance at a very competitive price. Then call StraightLine—the leader in 8mm and 4mm tape libraries. Now that's good thinking!

800-458-1273
www.igm.com

straightline

677 120th Ave. NE • Bellevue, Washington 98005
e-mail: igm@igm.com

© 1997 StraightLine. All rights reserved. StraightLine is a division of IGM Communications.
All brand and product names are property of their respective holders.

CIRCLE 166 ON READER SERVICE CARD

The first line is standard UNIX shell scripting and defines the interpreter that will be used to execute the script. The next three lines define the "submit" buttons used in the script. The fifth line identifies the *CGI.pm* library.

The sixth line instantiates a new CGI object that is manipulated via the *$editData* variable (you can use any variable name you like, assuming it's valid in Perl). The "constructor" for the object determines whether the form method is *GET* or *POST* and fetches the form data from the QUERY_STRING environment variable or STDIN respectively. The constructor does other nice things for you. It handles any necessary character conversions to "decode" the form data (converting "+" to space, etc.), and converts all of the "name=value" pairs into an associative array using "name" as a key for the "value." The seventh code line in the example checks for a "button" value. This will be a "null string" unless one of the "submit" buttons on the form has been clicked (more on this below).

The rest of the example determines which action to take during this invocation. If a field named "button" has the value "Submit Your Entry" and the *&editFormData()* call is successful, the script will process the data. In this case it formats and sends an e-mail message and then displays a successful result to the user, but you could just as easily save the data in a file or update a database.

If the *&editFormData()* call is *not* successful, the script will just fall back to the routine that displays the Web form to the user. And here is one of the best reasons to keep everything in a single script: when any edit fails, all the edit routine needs to do is print some error messages. These will then show up at the top of the form when it is redisplayed, and you will never again have to give instructions like "Field name is missing; press your browser's [Back] button and try again." When the form is repainted, the *CGI.pm* library retains any of the fields the user filled in and they will be redisplayed in the form.

If the form's "Help" button was used, the script can display the necessary help text to the user. If no "button" value was passed to the script, program logic falls through to the last "else" and the default "empty" Web form is presented to the user.

Self-referencing URLs

When you create your Web forms using a CGI script, you have control over how the form and any hyperlinks therein are created. One way to retain state within your application is to create recursive links that contain extra parameters. For example, suppose you have a submit form that creates a data file each time it is used. One function of your form could be to list some or all of the existing submittals on a page. You could add a link on each of the "open" requests allowing easy access to the data. Each line on the Web page could be created using something like the following:

```
$mySelf = $ENV{'SCRIPT_NAME'};
@openReqs  = &fetchOpenReqNumbers;
foreach $num (@openReqs) {
  print "<br> <a href=\"$mySelf?edit+$num\">
      Edit req $num</a>
  \n"; # end of print
}
```

The question mark separates the name of the script from the command-line parameters, and the plus sign separates subsequent parameters. This means that your form will need a mode that understands what to do with these command-line parameters, but it can simply be in another subroutine that reads the specified data file and prefills all of the fields in your form. (You already have a subroutine to display your form, *n'est-ce pas?*) At this point you can use your existing form to submit changes to the existing data.

The thing to remember is that your script will perform only one of its possible functions each time it is invoked. Once you are used to thinking about a single script as several separate operations, it becomes easier to envision how the subroutines can "overlap" a little. They can perform slightly differently depending on what you want to accomplish during each particular instance.

Command-line Parameters

When you pass command-line parms into your script as shown above, you need a way to figure out what they are. They are passed to your script just as you would normally expect in `$ARGV[0]`, `$ARGV[1]`, etc. What could be simpler? The trick is to remember what parameters you expect to see during the various instances of (or passes through) your script. One time your script builds a page with links like the ones above; the next time, when a link like this is used, your script responds accordingly. Note that the only time variables are available in `@ARGV` is when you use the "*?arg+arg*" syntax shown above. This will not be true when using the "*&name=value*" syntax shown below.

Prefilling Form Fields

Another handy thing that you can do is to prefill fields with data by passing "*name=value*" pairs through a URL. If you have a form with two fields named "Field1" and "Field2" and you invoke your script through a URL like this:

```
myServer/cgi-bin/myScript?Field1=FOO&Field2=LIKE+MAGIC
```

when the form is displayed, Field1 will contain "FOO" and Field2 will contain "LIKE MAGIC." The form is initially loaded with an HTTP *GET* request and, therefore, the parameters are available in the QUERY_STRING variable and *CGI.pm* interprets them correctly. In this case, the values will not appear in your script's command-line argument list.

You could also pass an initial "button" value into your form. Referring to the "main routine" example (see Form Buttons as Data) above, you could build a URL directly to the form's help text using:

```
<a href="/cgi-bin/myScript?button=Help+With+Submittal">
     Show myForm's help page</a>
```

Fields and Hidden Fields

Of course, a Web form's *raison d'être* is to collect input through the fields included on the form. A useful way to pass information between forms is through "hidden" fields. These can be thought of as a way to "cache" data in a form without it being easily changed by the user. The field data will be in the form, but no data entry field will be created. These are not truly hidden from the user, nor are their data secure from being changed by the user; however, they can be very convenient. For example, if a user starts filling out a form and then clicks the "Help" button, the script can build a form full of hidden fields on the Help page to retain the user's entries thus far. Then, when clicking an "Exit Help" button, the user's data is filled back into the original form. A quick way to do this using the *CGI.pm* method is the following example of the *&displayHelp()* subroutine. Assuming that our help page will be fairly long, let's put an "Exit Help" button at both the top and bottom of the page.

```
sub displayHelp {
    print " [ ... some heading stuff ... ] ";
    print $editData->startform;
    print $editData->submit('button',$Btn{'ExitHelp'});
    print "
      [ .. help text here ... ]
    \n";
    print $editData->submit('button',$Btn{'ExitHelp'});
    #
    # Now cache the user's data:
    #
    foreach $field ($editData->param) {
      print $editData->hidden(-name=>$field), "\n"
       if $field ne "button";
    }
    print $editData->endform;
    return(1);
}
```

You could also specify a different target window for your help page. This would make saving the fields unnecessary, but there may be times when you will need to save the data. To put your help page into a new window, change the "header" method to include the target window.

```
print $editData->header(-target=>'HelpWindow');
```

Another advantage of using a separate window for help: if you have several different places where the user can click the "Help" button, you don't have to figure out where the user was prior to clicking help (when clicking the "Exit Help" button, for example). If you use a separate window, be sure to remove any other navigation buttons from the help form so your users won't try to continue on within the Help window unless that is really what you want them to do.

Environment Variables

There are several handy shell environment variables created by the Web server, including:

- PATH_INFO — "extra" path information (not secure)
- SCRIPT_NAME — name of the running script (secure)
- REMOTE_USER — user ID when authentication is used
- REMOTE_HOST — may be a Web "proxy" server though
- HTTP_REFERER — may or may not be available

- HTTP_USER_AGENT reports the type of browser the user is running
- HTTP_COOKIE See "HTTP Cookies" below

Some of these can be considered secure and some of them cannot. PATH_INFO is one you cannot trust. If you add extra directory paths at the end of a CGI script, the extra part is stored here. Say you run a script named */cgi-bin/myScript* using the URL */cgi-bin/myScript/extra/path/info*. The Web server will search your CGI directory until it finds an executable and, in this case, put */extra/path/info* into the PATH_INFO variable. Never trust anything that users can send into your scripts.

The SCRIPT_NAME variable, however, can be considered secure information. The Web server will set this to the name of the script it invoked (relative to the Web Server's root directory). To continue the above example, it would contain */cgi-bin/myScript*.

HTTP_REFERER is another interesting variable; however, it may or may not be available. If your script was invoked from an "<a href>" anchor, this will contain that link. If your script was "opened" directly by typing it into the browser's location box, or opened from a user's bookmark, this variable will not exist.

Other variables may be available, depending on which server you are using. Some servers include all of the variables currently set by the user that starts the Web server running. Others, like Netscape's, do not, although you can add any by modifying the Netscape server's *obj.conf* file (and add variables such as HOME, PATH, etc.). See your server's documentation for details.

If you are not sure which variables are available or what their contents look like, build a little CGI script to dump them all out.

```
#!/usr/local/bin/perl
print "Content-type: text/html\n\n";
foreach $key (sort %ENV) {
  print "$key = $ENV{$key} <br> \n";
}
```

Remember that this will just show the environment as seen using the *GET* method. Several demo scripts accompany this article, and one called *getenv* allows you to display variables using either the *GET* or *POST* methods. See References below for details.

HTTP Cookies

Using HTTP "cookies" is a very cool way to cache data in the user's browser. Briefly, by adding some extra HTTP header information prior to starting an HTML document, you can pass data to the browser and then, during subsequent browses, request the information back. More recent versions of *CGI.pm* provide support for cookies and make them a snap. When you start your documents using the "header" method, just add the following:

```
print $editData->header(-cookie=>&cookieMonster);
```

Then, build a *&cookieMonster()* subroutine that will either store or fetch a cookie depending on which action is appropriate at the time. In this example, when the form's "submit" button is clicked, the values of two form fields will be stored in a cookie and, when someone using the same computer and login accesses this form again, these two fields will be prefilled using the values previously entered.

```
$cookieName="myCookie";
# Always look for an existing cookie first
%cookie = $editData->cookie($cookieName);

# We are storing field data, so pick which fields
@cookieFields = qw(SubmitterName SubmitterEmail);
if ($Button eq $Btn{'Submit'}) {
   # STORE submitter's info when they click Submit

  foreach $field (@cookieFields) {
      $cookie{$field}= $editData->param($field);
  }
} elsif ( ! $Button) {
  # FETCH previous cookie when starting a new form
  # and prefill the corresponding form fields

  foreach $field (@cookieFields) {
      $editData->param($field, $cookie{$field});
  }
}
# Now, bake a cookie that will stay on
# the user's machine for up to one year

$cookie = $editData->cookie(-name=>$cookieName,
                            -value=>\%cookie,
                           -expires=>'+1y');
return($cookie);
```

Each cookie you bake needs a name and this, combined with the machine and/or path on which it was created, uniquely identifies the cookie while it is stored on the user's computer. This means, of course, that the user must access your form from the same computer used the first time. There are limits on how many cookies you can create on one server and how much data each one can contain. A reference to more cookie dough (oops, make that cookie info) appears at the end of this article. Don't assume that your cookies will be there when you ask. There has been some negative press about the privacy issues surrounding cookies lately, and it is possible for users to configure their browser to deny acceptance of your cookies.

Data Files

The *CGI.pm* library provides an easy way to save the state of a form. The save method will even let you append data to an existing form file. This is one way to pass large amounts of data between multiple forms and it works well when combined with an HTTP cookie to keep track of a "user" or "session" ID tag.

```
open(DATA, ">/some/file");
$editData->save(DATA);
close(DATA);

open(IN, "/some/file");
$editData = new CGI(IN);
close(IN);
```

Of course, you can also access your own data files or a database from within your scripts if you'd rather not let *CGI.pm* do *all* your work for you.

Using Multiple Forms

(More than one logical form on one page.)

It is possible and often desirable to display multiple forms on one Web page. Just be aware that when a button is clicked, only the fields in that particular form (between the <form> ... </form> tags) will be returned to your script.

Formatting a page with multiple forms was difficult before tables came along, but now you have more control over the layout by simply putting a form in a table cell. You can have multiple buttons that each have a different "action" and still put them close together. Even if you have different subroutines that build each logical form on a single page, they can each use the *$var = new CGI* construct without problems. The *CGI.pm* library is designed to do the right thing.

An interesting use of this is a "fall through" selection form. Let's say that you have a Web application consisting of several different forms to store and display data based on some common selection criteria entered by the user. You could use a common subroutine to create the selection window(s) as one logical form at the top of each page, and each of the different form scripts could call the common routine to display the selection window above its respective "bottom" parts.

Take a look at *Figure 1*. The popup fields labelled "Select:" are in a separate form created by the select routine, and the bottom two "Edit Windows" are created by a second routine. Selections made in the top form will change the information displayed in the bottom form. The "Select" form also uses HTTP cookies and, when previous selections have been made, will simply return those selections. All other forms in this application that use the same selection routine act in the same manner, enabling users to navigate through their data without having to reselect anything until they are ready to do so. When they do, the select form saves the new values in the cookie.

When the selection form is called without any previous selections, it will wait until selections have been made. Once the "Enter Selections" button is clicked, the selection form will simply "fall through" and the bottom part will also then be displayed. This selection form is smart enough to wait when only one of the two possible fields has been entered (from a previous selection when only one of the choices was available). It will also wait for a new selection when one of the previously valid choices has been removed from the system since the form was last accessed. The bottom part of *Figure 1* has two logical forms—each button passes command-line parms back to the script through its unique "action" attribute.

Multipart Forms

(Forms that contain multiple items.)

A terrific recent addition to *CGI.pm* is support for forms created using "<multipart/form-data>" encoding. This allows you to use the "file upload" feature available in Netscape 2.0, and easily handles the necessary decoding that would be rather painful without a handy utility library.

```
print $editData->start_multipart_form;
print "<h2>Enter your name and a file to send</h2> \n";
print $editdData->textfield(-name=>SubmitterName,
```

```
            -size=>40, -maxlength=>55),"\n";

print $editData->filefield(-name=>FileName,
            -size=>40, -maxlength=>60);

print $editData->submit('button',$Btn{'SendFile'});
print $editData->endform;
```

The result of a "filefield" is a field in your form that looks like a "textfield" with a [Browse...] button immediately following the field. Clicking the button allows the user to browse files on the local computer and select one to upload to the Web server. When the user clicks a submit button, the contents of the file are transferred to the server along with data from other fields included in the form. The *CGI.pm* library separates the multiple parts of the form: all other form data is available as usual using the "param" method and, in addition, a temporary file is created. The temporary file will last only while your script is running, as *CGI.pm* will delete it just prior to exiting. If you want to keep it, you will have to copy it somewhere else. For example, part of your script's edit routine might include the following.

```
if ($Button eq $Btn{'SendFile'}) {
  $fileName = $editData->param('FileName');
  $tmpFile = $editData->tmpFileName($fileName);
  (@stat)  = stat($tmpFile);
  if ($stat[7] > 4096000) {
      &errorMsg("<b>$fileName</b> is too big (over 4MB)");
  } else {
      &copyUploadedFile($tmpFile);
  }
}
```

The *$fileName* is the name of the file that the user entered into the form, and the *$tmpFile* is the name of the temporary file created by the *CGI.pm* routines.

Of course, you don't have to put limits on how big a file your users can send to you, but if you don't, make sure you have enough disk space on your Web server to hold whatever people send. Note that if you do not use the "start_multipart_form" method, the file name will still be returned to your script but the contents will not have been uploaded to your server. This can actually be a good side effect. Using a "filefield" in a form created with the "$editData->startform" method is nice if the user will just be identifying files that already exist on the Web server. In order to point to them as "local" files the user would need to be running the browser on the Web server (or need to have access to the files via an NFS mount or a network "share" drive) but, if this is the case, you can avoid uploading files when they're already there.

Multipage Forms

(Forms presented on more than one "page.")

One of the nicest things about Web forms is that HTML automatically builds scroll bars, allowing you to concentrate on content instead of the implementation. You can now build longer forms than you ever thought possible; however, there will be times when you need to break up the form into multiple "pages."

As you add more complexity to a form, the number of possible fields may get too cumbersome to handle all on one "page." Collecting some data on a first page and then building a second page based on the first can make a form more usable. It is possible to build a form so complex that a browser simply cannot display it entirely or, if it does display, scrolling and data entry can be so slow that the form is unusable. When this happens, you may need to use "hidden" fields and/or data files to pass information on to another form in your application.

Multimode Forms

(Forms that perform more than one function.)

Now that you are using a CGI script to create your Web form, you can build variations of the form at different times depending on who is running the script. For example, let's say that you have a form to submit a request, and the request must be approved by an administrator, editor, or someone in charge of authorizing the request. You could build an entirely separate version of your form, or you could add an "admin" mode to the script you already have. The trick here is to differentiate when the script is run by one of the application administrators and when it is run by someone submitting a request.

One way to accomplish this is to create an "adm" subdirectory underneath your application's CGI directory and add Web security to the scripts in this subdirectory. In the subdirectory simply create symbolic links to the actual scripts and, when they are run from this subdirectory, the user will be prompted for a user ID and a password. To make this work, your Web server must be configured to follow symbolic links. For example:

User access: `<a href="/cgi-bin/myApp/myScript">`
Admin access: `<a href="/cgi-bin/myApp/adm/myScript">`

The above example assumes the following symbolic link exists.

```
cgi-bin/myapp/adm/myScript -> ../myScript
```

The previous example also assumes the subdirectory *cgi-bin/myApp/adm* has basic Web authentication in place (consisting of an *.htaccass* or an *.nsaccess* file).

You can then check for a REMOTE_USER environment variable, or you can check for the string */adm/* in the SCRIPT_NAME variable. Once your script determines that it is in "admin" mode, you can then alter your form accordingly. For example, you can add "approve" and "deny" buttons, add a text field for "admin comments," etc. You could even change the default "screen" for admin users to be a list of all currently outstanding requests where each entry in the list was a recursive "" link to your same script. Of course, this means that you need to add an "edit" mode that can handle this type of access. However, your script is building the list and you have complete control over how you choose to build the links and act on them.

Secure Scripts

The above method is a simple way to provide an "admin" login; however, you may want everyone using your particular application to log in using a Web password so you can always tell who is doing what with your forms. In this case you may want to consider adding user "classes" to your application. This can be done quite easily right on top of the existing Web security mechanism simply by rethinking how you use the Web authentication "group" file.

Don't forget that using your Web server's "basic" security is just that—basic. You should not assume that your application is protected just because you have added Web security. Within Hewlett-Packard, for example, company confidential information can be placed on the intranet (internal network) with basic Web security; however, this author would not choose to put that same information on the open Internet without using a better security scheme, such as SSL.

Building on Basic Web Security

If you place basic Web security on your CGI directory or the CGI subdirectory where your application's scripts are located, everyone who accesses your scripts will be prompted for a user ID and password. After this, you can determine who is running the script via the REMOTE_USER variable. It is also possible, as part of basic security, to create a "group" access file and require that people be a member of some group to access your application. However, the group information is not then available to your scripts. As an alternative, simply restrict access to a "valid user" and allow anyone with a current user ID and password access. Do still maintain a "group" file though, and add users to one or more groups as appropriate. It is very easy then to parse your group file and build a list of those groups of which the current user is a member.

At this point you can build user "classes" into your Web scripts. One way, for example, would be to define three classes: "admin," "user," and "guest." If the user is a member of the "guest" group, then he has only guest access; if a member of the "admin" group, he has admin access. If neither of these applies, then he is considered a "vanilla" user.

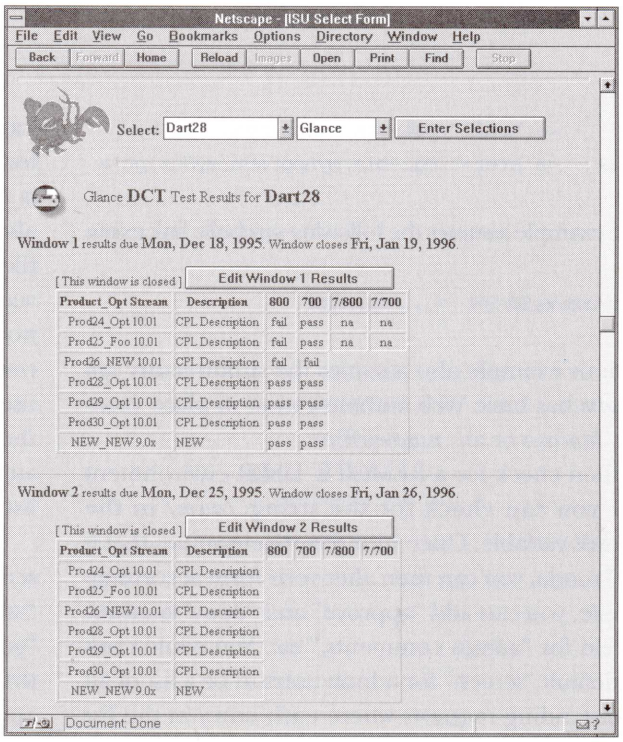
Remember that your script will need to parse for group access each time a button is clicked. An alternative here is to cache the group list in a cookie that lasts only until the user exits the browser program. This is not completely safe, as sophisticated users could spoof your application by sending their own cookie strings to your scripts. If you do use this method, be sure that the cookie is not retained on the client machine. Omit the "expiry" date on a cookie, and it will exist only until the user exits the browser. You don't want to make it *too* easy for your users to modify their own access capabilities.

You could further subdivide the user classes:

- user ID "admin" in group "admin" has greater privilege than other users in group "admin"
- user ID "guest" in group "guest" has less privilege than other users in group "guest"
- create "group" administrators using user ID "fooadmin" in group "foo" (or in group "admin"), or create a group named "admin-foo" into which you add users who will have admin access within group "foo"

The possibilities are limited only by your imagination, the time you have to spend, and the requirements of your particular application. After this, however, it will be up to your scripts to determine if someone has access to a given part of your application, and how your scripts will react when unauthorized access is requested.

FIGURE 1



You may want to design your scripts to filter out all choices, buttons, etc., that are unavailable at each security level. Only showing your users what they can do and omitting what they cannot will help eliminate their confusion and frustration. The selection form at the top of *Figure 1*, for example, will only contain two popup menus when the user has been added to more than one "group" (in this case, Glance and at least one other). Otherwise, the user would only see the "Release" popup menu (in this case, Dart28, etc.). In the bottom section of *Figure 1*, the "edit window" is closed and non-admin users would not have an "Edit" button to click. The user in this example is an administrator, so they can edit the data.

Adding a Registration Form

One nice thing to do at this point is modify your Web server's configuration to alter how "401" errors are handled. The details vary depending on your server. For NCSA's httpd you can add an "ErrorDocument" directive into the *srm.conf* file. For Netscape's products, you can either use their handy admin forms to enable this or you can add an "Error reason=" entry directly into the *obj.conf* file (refer to your Web server's documentation for specifics).

Then, when access is denied, the server runs a script of your design to display an alternate error message, and the script can provide a link to a registration form allowing users to request access to the private parts of your applications. Since doing this changes the action of *all* "401" errors on your Web site, you may want to make this script smart enough to know when it should display a link to the registration form and when it should simply display a default error message. Examine the HTTP_REFERER and/or PATH_INFO environment variables and proceed based on where the "401" error occurred. Unfortunately, depending on which Web server you are running, you may need to use other variables to get this information.

You then have the best of both worlds: any 401 errors generated by your Web server will be handled automatically, and your scripts can call your error handler when appropriate. The following example assumes that your error handler script is named "accessDenied," and your application script has determined the user does not have access to the requested information.

```
print "Location: /cgi-bin/accessDenied\n\n";
```

Be sure that you have not printed anything else to STDOUT and that you do include the two carriage returns, as this is what separates the HTTP "header" information from the rest of an HTML document. Exit your script soon after this, as nothing else you print will make it to the user's browser.

One special note: the "Location:" header has officially been replaced by the "Uri:" header. A nice thing about the "Location" header is that using a relative URL does not change what the user sees in the browser's location box, while using a complete URL will change the location shown in the browser. In the above example, the user will never know that the "accessDenied" script was run. Using *CGI.pm*'s redirect method will print both the unofficial "Location" header and the official "URI" header; however, it expects a completely qualified URL. This may not be what you want to do.

```
print $editData->redirect('http://myserver/cgi-bin/accessDenied');
```

Either of these last two examples will immediately direct the user's browser to whatever URL you specify. What does the "accessDenied" script look like? Let's assume that your admin scripts are located in the *cgi-bin/adm* subdirectory and are protected with basic Web authentication. All you need is the following.

```
#!/usr/local/bin/perl
print "Content-type: text/html\n\n";
print "<HTML><HEAD><TITLE>Unauthorized</TITLE></HEAD>\n";
if ($ENV{'PATH_INFO'} =~ m#^/cgi-bin/adm/#) {
  print "
  <BODY><H2>Access Denied</H2>
  Authorization is required to access this application.
  <H2>Registration is Available</H2>
  If you would like, you can
  <a href=\"/cgi-bin/newUser\">register for access</a>
  to the Administrative Forms
  \n";
} else {
 ## Default Netscape Server Unauthorized response:
  print "
  <BODY><H1>Unauthorized</H1>
  Proper authorization is required for this area. Either
  your browser does not perform authorization, or your
  authorization has failed.
 \n";
}
print "</BODY></HTML>\n";
exit(0);
```

Remember that the "PATH_INFO" variable is not secure. So, why do we use it here? Different Web servers handle error redirection differently. The Netscape Enterprise Server sets this variable instead of the "SCRIPT-NAME" when it runs your custom error handler script. The NCSA Web server uses the "REDIRECT_REQUEST" and "REDIRECT_URL" variables. You will probably have to add "print" statements to display the available variables until your script is working just right. See the example above in the section on "Environment Variables" for one way to do this.

How you choose to design the "newUser" registration script/form is entirely up to you. Just be sure it is not in a protected subdirectory or no one will be able to register.

Template Forms

(Forms that can be prefilled and saved on a local computer.)

As more and more processes migrate to the Web, there are times when it would be nice to allow your submitters the ability to prefill a form and save it to their local disk drive as a "template." The user can use the "File" and "Save as..." menus on the browser to save a copy of the form locally, and then later use the browser's "File" and "Load file..." menus to recall the template. This is easily accomplished, but there are four things necessary to ensure success.

First, be sure to use a complete URL in the form's "action" tag, including the machine name and port number (if other than 80). Otherwise, when the user recalls the template, the browser will be unable to submit the form back to the Web server. (You could also choose to add a "<base>" tag to your form to resolve any relative links, and *CGI.pm*'s "start_html" method supports a "base" parameter.)

Second, the file that is saved to the user's local disk will not contain any of the user's input unless a "submit" button was clicked to alter the form prior to being saved. In other words, entering data, selecting from popup menus, etc., does not change the form's "source." You can verify this by filling out the form and then clicking the "View" and "Document source" menus. To make this work, add a "Make Template" button to your form that will simply redisplay the form with the user's input. As this does change the form's "source," the changes will be included when the user saves the file. In addition, you can choose to skip any or all of the form's normal edits, and empty out any date fields that should not be saved.

Third, if you have icons or other graphics on your form, be sure to fully qualify the URLs to the images or they will appear as broken graphics when the user reloads a saved template from the local disk. Adding a "<base>" tag to your form, as mentioned above, will also solve the problem of images with relative links appearing broken later on.

Fourth, add a version number in a hidden field. If it becomes necessary to modify the form in such a way that previously saved templates are rendered unusable, you can notify the users when they attempt to submit an out-of-date form. This is not foolproof since, of course, the users could simply edit their template file and modify the version saved in their template, but it should usually be sufficient.

At this point, since you are recreating the form, you could add special instructions to the top of the template with reminders

on usage. These will still be at the top of the form when it is loaded from disk, so any instructions must make sense both when saving the template and when retrieving it later.

You may want to determine how many people submit via saved templates. When they first click a "submit" button, the HTTP_REFERER environment variable will contain "file:" as the protocol part of the URL. This will only be true the first time. If edits fail, for example, or a "help" button is clicked first, any subsequent submittals will no longer be from their file but from your Web server. You will have only one chance to determine if a template was used, so grab this information and tuck it away in a hidden field.

Potential Technical Problems

The biggest problem with CGI is security. It opens a little window into your system that is only as secure or insecure as you make it. First off, be sure to run your Web server by some user other than "root" or, if you must do this, use a server that can run your scripts as a non-root user. This way, if anyone is able to compromise your scripts, they won't have "super user" access to your computer.

Next, learn as much as you can about your Web server's security configuration files. You should **not** configure the following settings without good reason:

- allowing "server-side includes," (AKA: "parsed" HTML), especially with "CGI execs"
- enabling ".cgi" as a "mime" type, allowing anyone to put a CGI script anywhere
- allowing anyone to override your server's configurations in their local directories
- giving everyone "write" permission to your CGI directory

Keep all of your CGI programs in one directory tree underneath your server's "root" directory, and make sure anyone who puts scripts on your server understands the risks. The simpler the setup, the easier it is to keep secure.

Never Trust Input Data

Once again, never *ever* trust form data in your CGI scripts. It is far too easy for anyone to submit bogus data through your form even if you use "hidden" form fields. Perl has a command line option to help you protect yourself from trouble. Get familiar with using the "-T" option on your scripts. This tells Perl to flag any variables with user data as tainted. Perl will then prevent you from using these variables in insecure ways (file opens, system calls, etc.). You must use Perl's match functions to "untaint" the variables and, while this still is not completely secure, it forces you to think about potential problems in advance.

Design Around Missing Utilities

If your script moves from one system to another, any external programs you run may not exist. Even OS utilities may not be there—sendmail moves when you move from HP-UX 9.x to 10.x. Be sure to verify that you can run external programs.

```
if (-x "/usr/lib/sendmail") {
  $mailer = "/usr/lib/sendmail"; # 9.x
} elsif (-x /usr/sbin/sendmail) {
  $mailer = "/usr/sbin/sendmail"; # 10.x
} else {
  &systemErrorMsg("Can't run sendmail");
}
```

What happens in our "main routine" example if the *&sendEmail()* routine fails even with the error checking shown here? There may not be enough process space left to run another program, and other problems can occur. Just create a slightly different version of your *data* error routine for *system* errors. Have it display an appropriate message and abort your script. Assuming the problem is not with sendmail, the error routine could even send you e-mail about the problem.

Part II: Graphic Design Issues

Form Layout

Creating Web forms is really Graphic User Interface design in disguise. It is appropriate to spend a considerable amount of time ensuring a consistent look and feel across all of the Web forms in one, some, or all of your applications. One layout that is particularly flexible to use is a four-column grid.

```
Submitter Name [    ]   Submitter E-mail [    ]
```

This works exceptionally well when combined with a two-table layout. Use one "outer" HTML table simply to create a border of one or two pixels around your whole form. Create a second "inner" HTML table to do all of the column formatting from the top to the bottom of your form. Specify a

border of "0" for your inner table and, when you want to create visual groupings within your form, use a hairline rule that spans the four columns.

```
print "
    <table border=2 width=100%>
    <tr><td>

    <table border=0 width=100%>
    <tr>
          <td align=right> Submitter Name  </td>
          <td align=left>",
                $editData->textfield(-name=>"SubmitterName",
                      -size=>20, -maxlength=>35),"
          </td>

          <td align=right> Submitter E-mail  </td>
          <td align=left>",
                $editData->popup_menu(-name=>"SubmitterEmail",
                      -size=>20, -maxlength=>45),"
          </td>
       </tr>

       <!---------------------------------- >
       <tr> <td colspan=4>  <hr>  </td>  </tr>
          [ ... rest of form here ... ]
       </table>

       </td></tr>
       </table>
\n"; # end of print statement
```

If you start with this layout, all of the fields will be nicely aligned. You can use the "colspan" and/or "rowspan" attributes to combine multiple columns and/or rows to obtain precisely the desired effect. Think of this as similar to a layout "grid" used in desktop publishing. It has many of the same advantages. Not all of your forms will work in this format, but you may find that many work nicely.

Create Meaningful Error Messages

A good thing to do for your users is create clear errors, and a good way to do this starts with using meaningful field names in your form. Instead of using "Field1" and "Field2" or even "Name" and "Email" consider using the same wording for field names that you use to label the fields on your form. (See the example just above in "Form Layout.") This way, not only will your scripts be easier to maintain, but you are setting yourself up for easy error handling.

Remember the very first example in this article where the overall program flow was outlined? (See "Using Form Buttons as Data" above.) When the "Submit Your Entry" button is clicked, the edits we want might look something like this.

```
sub editFormData {
  $Ok = 1; $NotOk = 0;
  @requiredFields = qw( SubmitterName SubmitterEmail );

  $errorFields = $separator = "";
  $errorNums=0;

  foreach $field (@requiredFields) {
      if (! $editData->param($field)) {
        $errorFields .= $separator . $field;
        $errorNums++;
        $separator = ", ";
      }
  }
  if ($errorFields) {
      &errorMsg("The required field(s) are empty:
                  <b>$errorFields</b>");
  }
  #
  # Other edits go here...
  #
  if ($errorHeadingsPrinted) {
      &errorMsg("EOD");
      $editErrorsOccured="True";
      return($NotOk);
  }
  return($Ok);
}
```

The "@requiredFields" variable is used for two reasons. First, since the field names and their values are stored by *CGI.pm* in an associative array, they are in no particular order. This way, when any required fields are empty, the error message will display the list of missing fields in the same order as they appear on the form (assuming that you are careful when you build this list). Second, you may want to have a different

list of required fields, depending on what button was clicked or what action you are preparing to take next. In this case, it is easy to add a test and build the list as necessary.

What does the *errorMsg()* routine look like? Grab the example scripts available on the Web (see the URL in the byline at the bottom of this article). Briefly, it prints a heading and the error message the first time through and, on each subsequent call, just prints the error message. After all your edits have been performed, check to see if any errors have occurred. If so, call the routine one last time with "EOD" so it will finish up the error heading block. Just add an "<hr size=4 noshade>" tag to make the errors more obvious and a bit of generic wording to "make changes and click [Submit] again". See *Figure 2.*

To make the error messages even more obvious, you can choose to omit the form's regular heading as shown in *Figure 2.* This works very well. When the main routine determines that edit errors occurred, it simply falls back through and displays the form again. All the errors will appear at the top, and the *CGI.pm* library retains any of the fields that were previously entered. Once the necessary fields are entered and all edits pass, the script continues to process the request and display the confirmation page.

Counters

Our development team won't use access counters in any of our applications—they are mostly irrelevant, usually misleading, and require far too much overhead to create and download. We prefer to generate meaningful Web statistics reports and make them available to those who are interested in the information.

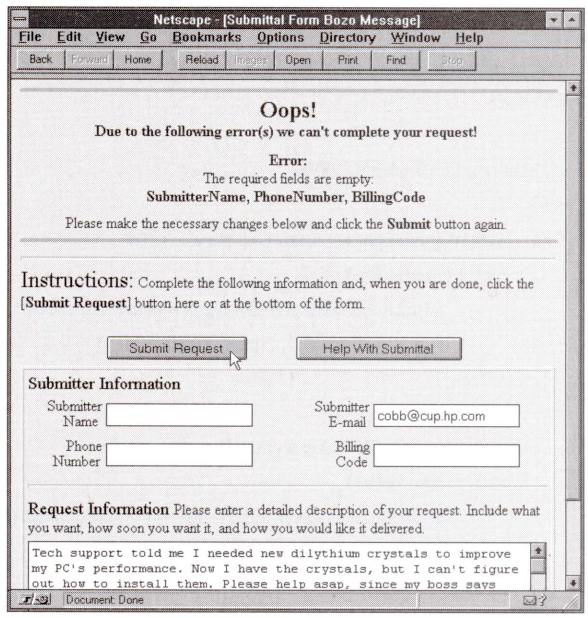
Potential Design Problems

You don't have to become a graphic designer to build good-looking, usable Web forms, but gaining a basic understanding of graphic and user interface design issues will certainly help. If people find your forms difficult to understand, they will tend to avoid your Web site altogether.

Design for Forward Motion

Relying on a browser's "Back" button is a good way to confuse your users. Designing your forms with an obvious button or link to the next logical step will help prevent people from going back to forms with stale data. If they have just submitted a "delete" transaction through your form and then click the "Back" button, they may not realize the deletion was success-

FIGURE 2



References

The Web, of course. Complete examples of these concepts are available from the URL in the byline. Documentation for the *CGI.pm* library, including installation instructions, is available from Lincoln Stein, director of informatics, Whitehead Institute Center for Genome Research at the Whitehead Institute for Biomedical Research. For instructions, see "*http://wwwgenome.wi.mit.edu/ftp/pub/software/WWW/cgi_docs.html*".

Dr. Stein also includes good example scripts with his libraries, and he has other Perl tools available too.

For further information on HTTP cookies, visit Netscape at "*http://home.netscape.com/newsref/std/cookie_spec.html*".

You should also check out CPAN, the Comprehensive Perl Archive, where you may find just the tool you've been looking for (or you may have a cool routine to contribute). To find the mirror site closest to you, use your favorite search engine and look up "CPAN."

ful and attempt to submit the same transaction once again. This will probably result in an error, assuming the original deletion was successful, and further confuse the situation.

Design Around Missing Graphics

Just as with the external programs your script runs, do not assume that your logos and other graphics will be available. Those broken icons are so annoying, and they can be avoided so easily.

```
$docRoot="/some/path/to/htdocs";
$imageFile="/images/myLogo.jpg";
if (-r "$docRoot/$imageFile") {
  print "<img src=\"$imageFile\"> \n";
} else {
  print "<h2>My 'Lil Form</h2> \n";
}
```

Use the ALT Attribute

Not everyone has fast access to the Net, especially these days! You may want to add text to your forms that is used in place of graphics when users choose not to download graphics as they browse the Web.

```
<h2> <img src="myLogo.jpg" alt="Submit Form"> </h2>
```

If you wrap the image tag in a heading tag, the alternate text will be displayed as a heading. You may want to do this in other places throughout your Web site. This way, people who do not automatically download graphic images because of network limitations can get a similar "look and feel" from your Web pages.

Summary

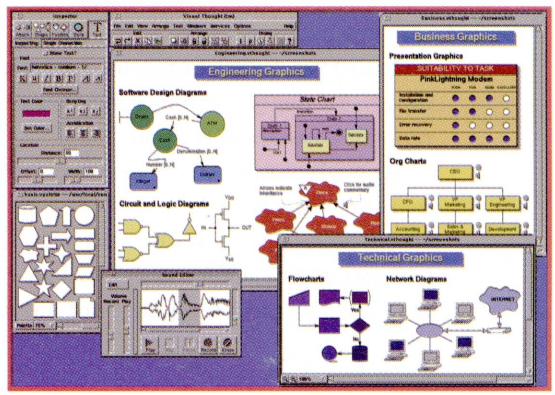
There are more tips and tricks to Web forms than can fit into this article. Some other techniques you may want to investigate include non-parsed header (NPH) scripts, JavaScript, and, if your Web server supports it, FastCGI. The *CGI.pm* library has, or soon will have, support for these in addition to all the other clever things it does. The Web evolves at the speed of light, and the tools you embrace should not hold you back.

I hope you have enjoyed this brief look at some of the CGI techniques we've been using and adapting in our Web projects. If most of this was familiar to you, I hope that you at least found something of interest. If this was mostly new to you, once you start using this, add "Object-oriented, event-driven Web application designer" to your résumé. Then get out there and wow 'em! ◧

Chris Cobb is a senior engineer with HP Worldwide Sales Communication within the Computer Systems Organization in Cupertino, California. His e-mail address is cobb@cup.hp.com and his Web URL is http://www.cup.hp.com/~cobb.

visual thought®

UNIX Diagramming & Flowcharting



Use Visual Thought for:
- Software design diagrams (Booch, Rumbaugh, Objectory, Fusion, custom notations)
- Flowcharts for ISO 9000, TQM, BPR
- Clickable Web diagrams with GIF/JPEG and server/client-side imagemap export
- Framemaker documentation graphics w/ MIF, EPSI export
- Network diagrams with network component clip art
- Dataflow diagrams, org charts
- Presentation & business graphics
- Graphical pre/postprocessor for other tools or simulators

FREE CD-ROM, FTP & 30-day trial

Additional benefits
- Intelligent, rubberbanding lines
- Macintosh-like ease-of-use . . . on UNIX!
- Editable, WYSIWYG drag-and-drop palettes
- Hyperlinked, hierarchical documents
- Arbitrary object rotation
- Complete text handling: subscripts, arbitrary fonts, sizes, colors, styles and justification
- Dozens of export formats: GIF89, JPEG, MIF, TIFF, EPSI, XWD, SunRaster, others
- Available on SunOS, Solaris, HP-UX, and soon, Windows 95/NT

Confluent®

Confluent, Inc.
132 Encline Court
San Francisco,
California 94127
415-586-8700 Voice
415-586-8838 Fax

See why Visual Thought has thousands of users worldwide! Clients include: ABB • AT&T • Alcatel • BNR • Bellcore • Boeing • CS First Boston • Credit Suisse • E-Systems • EDS • Ericsson • Ford • HP • IONA • Lucent • McDonnell Douglas • Motorola • Nokia • Nomura • Nortel • Persistence • Prediction • Pure Software • SAIC • Smith Barney • Sprint • Toshiba • U.S. Air Force • U.S. Navy • US West • UBS • Wells Fargo Nikko

VISIT (use code AE14)
http://www.confluent.com/
800-780-2838 ext.154
info@confluent.com

©1996 Confluent, Inc. All rights reserved. Visual Thought, the Visual Thought logo, Confluent, and the Confluent logo are trademarks or registered trademarks of Confluent, Inc. All other names are the properties of their respective holders.

CIRCLE 37 ON READER SERVICE CARD

HP-UX

by David Totsch

How Are Things Going? Part II

LAST TIME, WE EXPLORED the automation of file system space monitoring as a beginning to a series of columns dedicated to training HP-UX to monitor its own status and report on an exception basis.

If you deployed a file system monitoring script, the first time it reported a file system over threshold you probably asked yourself something like, "OK, I have a file system filling up—now what do I do?" The wonderful thing about UNIX is that there is almost never a single solution to any one problem. You could use any of the following approaches (and a few I have not thought of and several I cannot put into print):

- extend the logical volume
- introduce a new file system to hold some of the data
- compress some files
- migrate files to backup media
- clean up log files

Once you have finished entertaining visions of physical and/or mental anguish for certain users, you will probably opt for activities that are productive and not disruptive to users (not to mention your personal freedom). Provided the file system in question has them, examination and truncation of log files can be very productive. Clearing log files is non-disruptive and can be very effective at recapturing disk space.

First of all, we have to have a list of log files to examine. *Listing 1* is by no means a comprehensive list, but it is a starting point. When you build the list for your system, remember to include the log files created by third-party software.

Each file needs to be checked against a threshold. If it is over the threshold, one of the following actions can be taken:

- simply report it
- truncate the file
- rotate the log file

A log can be selected for action based upon one of at least three factors: characters, lines, and blocks. Only reporting the log file still leaves you with the burden of a manual task. Rotating the log into a copy is not complicated, but be sure that the log file is not tracked by the application through the inode number. A very simple approach is to truncate the log file in question. If you know the current number of lines in the log and how many you want to limit it to, you can use the following Korn shell fragment to truncate:

```
# LOG    - the log file in question
# CLINES - current number of lines in the log
# MLINES - maximum number of lines in the log
ed - ${LOG} <<- =EOI=
        1,$(( ${CLINES} - ${MLINES} )) d
        w
        q
        =EOI=
```

LISTING 1

```
/etc/eisa/config.log
/etc/shutdownlog
/var/adm/automount.log
/var/adm/backuplog
/var/adm/btmp
/var/adm/cron/log
/var/adm/diag/*MEMLOG
/var/adm/diag/LOG*
/var/adm/lp/log
/var/adm/lp/log.piftopcl
/var/adm/lp/lpana.log
/var/adm/lp/lpd.log
/var/adm/lp/lpdaemon.log
/var/adm/messages
/var/adm/msgbuf
/var/adm/ptydaemonlog
/var/adm/rbootd.log
/var/adm/rld.log
/var/adm/rpc.lockd.log
/var/adm/rpc.statd.log
/var/adm/snmpd.log
/var/adm/spellhist
/var/adm/sulog
/var/adm/sw/install.log
/var/adm/sw/swagent.log
/var/adm/sw/swagentd.log
/var/adm/sw/swconfig.log
/var/adm/sw/swcopy.log
/var/adm/sw/swinstall.log
/var/adm/sw/swmodify.log
/var/adm/sw/swpackage.log
/var/adm/sw/swremove.log
/var/adm/sw/swverify.log
/var/adm/syslog/mail.log
/var/adm/syslog/syslog.log
/var/adm/vtdaemonlog
/var/opt/dce/config/dce_config.log
/var/opt/dce/rpc/rpcd.log
/var/sam/log/br_log
/var/sam/log/samlog
/var/tmp/sam_remove.log
/var/tmp/swagent.log
/var/uucp/.Admin/DIALLOG
/var/uucp/.Admin/Foreign
/var/uucp/.Admin/audit
/var/uucp/.Admin/errors
/var/uucp/.Admin/uucleanup
/var/uucp/.Admin/xferstats
/var/uucp/.Log/LOGX25
/var/uucp/.Log/uucico/*
/var/uucp/.Log/uucp/*
/var/uucp/.Log/uux/*
/var/uucp/.Log/uuxqt/*
/var/yp/ypserv.log
/var/yp/ypxfr.log
```

LISTING 2

```
# Awk script to
scan su log...
# Establish number of failures per day before reporting:
maxfail=3    # one user su to another
maxsufail=5  # anyone su to root
$4 ~ /-/     { # on every failed su..
      split($6,ids,"-")
      # ids[1] is user running su
      # ids[2] is the user they wanted to become
      data[$2","ids[1]","ids[2]]++
      # $2 is the day date notation; increment the
      # the count for this user su to ids[2]
      if ( ids[2] == "root" ) shots[$2]++
# keep track of total failures to root, just in case...
      }
END   { # report...
      for ( x in data ) if ( data[2] > maxfail )
            {
            split(x,rpt,",")
print
"On",rpt[1],rpt[2],"attempted",rpt[3],data[x],"times." }
      for ( y in shots ) if (shots[y] > maxsufails)
            {
print "On",y,"there were",shots[y],"failed attempts at
root!" }
}
```

There is one limitation you should be aware of: you need enough space in */tmp* to hold the log file you are truncating. That can be determined by comparing the output of *du(1)* to the output of *bdf(1)*.

When a file is truncated, information is thrown into the bit bucket. If the task is automated, you may be tossing important information that you never had a chance to evaluate. Think of *sulog*. You would not want to discard any contents that contained excessive failed attempts. *Listing 2* is a very simple *awk(1)* script to report those failed attempts. It checks for excessive failed attempts by one user during a particular day. It reports on the entire contents of the file and has no mechanism to prevent multiple reporting of the same violations.

Unfortunately, each set of related log files will require its own automated examination routine. That has a severe implication for log files that store their information in non-ASCII format. Data log files have to be formatted into an ASCII format and reported on; their data removed selectively, they must then be returned to their data format. Next time, we will explore some techniques that can be used to format and report on two binary log files that can grow without bounds: *btmp* and *wtmp*. ◧

David L. Totsch is a technical consultant for Premier Systems Integrators, Inc. in Charlotte, North Carolina. His specialty is HP-UX system administration and he enjoys training others to do the same. He can be reached at 704/522-6088 or totsch@rbdc.rbdc.com.

NT Process Scheduling

by Bob Combs

WE'VE ALL CODED PROGRAMS that schedule other programs. In fact, it's hard to build any kind of sophisticated application that doesn't rely on programs cooperating with one another. Even client-server processes schedule other processes. A client program may schedule a search request form or a dial-up networking program. A server process may schedule a background batch process or a message-specific program. Well, you get the picture ... the list goes on and on. Of course some of these processes could be scheduled as threads of the client or server program.

The initial process that schedules another proce I'll refer to as the parent process or program. The scheduled process I'll refer to as the child process or program.

Under HP-UX (or UNIX in general, for that matter) a parent program schedules a child program using an amoebic mechanism. The parent process first clones a version of itself with a `fork` call, and then the child process (the cloned one) makes an `exec` call to overlay itself with the desired program. The child process then begins execution of the scheduled child program. It is the `fork` call that determines whether it is returning from the parent or child process by the return value. The `fork` call returns a 0 to the child process scheduled, and the process ID (PID) of the child process to the parent. The code for this operation (which we've all seen a thousand times) is noted in *Listing 1*.

Microsoft Windows NT has a POSIX subsystem that is POSIX.1 compatible and supports the `fork` and `exec` function calls. (The `exec` call has been available in Microsoft's C run-time library forever.) However, there are a few catches. POSIX applications expect only the NTFS file system. Programs can be scheduled from any NT file system, but can access only an NTFS file system. POSIX file calls operate a little differently than similar calls for non-POSIX file calls. The user right *`Bypass Traverse Checking`* must be turned off for true POSIX compliance. All POSIX applications execute in a command prompt window since they don't support the Windows GUI. The POSIX system must be disabled if C2 security is required.

As you can see, POSIX applications may not interact correctly with other non-POSIX applications. And Microsoft discourages the use of the POSIX subsystem. It's really there only to allow quick ports of UNIX applications to Windows NT. The preferred and compatible method is to code the application native to Windows NT with the Win32 API set. Of course, Win32 uses a slightly different paradigm.

Under Win32, the parent process schedules the child process with one API call; no process cloning or overlaying is involved. The child process is started while the parent continues to execute. The Win32 call to schedule a process is *`CreateProcess`* as described in *Listing 2*. While this mechanism is more straightforward than the UNIX `exec` call, the calling parameters are more complex.

The first two parameters of the *`CreateProcess`* function describe the file name of the program to be scheduled and the parameters to be passed. The *`lpApplicationName`* parameter is the file name to schedule, unless the application is a 16-bit DOS or Windows 3.1 program. In that case, *`lpApplicationName`* should be NULL and the file name is given as the first token on the *`lpCommandLine`* para-

meter. The file name to schedule may be in either parameter if it is a 32-bit program. Therefore, as a general case, I make *lpApplicationName* NULL and place the program file name and parameters in *lpCommandLine*. Note in *Listing 3* that I set the second parameter to MyProg and have no parameters.

The *lpProcessAttributes* points to a security attributes structure that describes the security context to be given to the child process object. Similarly *lpThreadAttributes* points to the security for the child primary thread object. Normally I set both these parameters to NULL, requesting the default security descriptors.

I've set the *dwCreationFlags* to *CREATE_SEPARATE_WOW_VDM*, which is ignored if it's a 32-bit application and otherwise creates the child process in a separate Windows 3.1-compatible (Windows On Windows) virtual DOS machine (VDM) process space. Normally Windows 3.1-compatible processes are created as another thread in one system-wide WOW VDM. The single system-wide Windows 3.1 VDM more closely simulates how Windows 3.1 works. I prefer to run each 16-bit application in a separate VDM so it cannot interfere with, or be interfered with by, other 16-bit applications.

Since *CreateProcess* expects the caller to supply startup information, I call *GetStartupInfo* to read the *STARTUPINFO* structure (*Listing 4*) of the parent process. The structure specifies how the new window should appear. Upon successful completion, *CreateProcess* will return in *lpProcessInformation* a *PROCESS_INFORMATION* structure (*Listing 5*) that describes the child process created.

So far, the parent process has started the child process and could at this point resume executing while the child executes. But what if we want to wait for the child to complete its task and then pick up the return code from the child? Then we need two additional functions.

The *WaitForSingleObject* Win32 API will wait for the

LISTING 1 *UNIX Scheduling Example*

```
switch(fork())
     {
          case -1:
               printf("fork failed.\n");
               return;     /* error return */
          case 0:
               execl("someProgram", arg0, NULL);
               /* child success won't get here */
               printf("exec failed.\n");
               exit(1);
          default:
               if(wait(NULL) == -1)
                    printf("wait failed.\n");
               /* parent success */
     }
```

LISTING 2 *CreateProcess Call*

```
BOOL CreateProcess(
    LPCTSTR lpApplicationName,                // pointer to name of executable module
    LPTSTR lpCommandLine,                     // pointer to command line string
    LPSECURITY_ATTRIBUTES lpProcessAttributes, // pointer to process security attributes
    LPSECURITY_ATTRIBUTES lpThreadAttributes,  // pointer to thread security attributes
    BOOL bInheritHandles,                     // handle inheritance flag
    DWORD dwCreationFlags,                    // creation flags
    LPVOID lpEnvironment,                     // pointer to new environment block
    LPCTSTR lpCurrentDirectory,               // pointer to current directory name
    LPSTARTUPINFO lpStartupInfo,              // pointer to STARTUPINFO
    LPPROCESS_INFORMATION lpProcessInformation // pointer to PROCESS_INFORMATION
);
```

LISTING 3 *Win32 Scheduling Example*

```
DWORD dwErr, dwValue;
PROCESS_INFORMATION pi;
STARTUPINFO sui;
GetStartupInfo(&sui);
if(!CreateProcess(NULL,
                  "MyProg",
                  NULL,
                  NULL,
                  FALSE,
                  CREATE_SEPARATE_WOW_VDM,
                  NULL,
                  NULL,
                  &sui,
                  &pi)) {
     dwErr=GetLastError();
     printf("Unable to schedule MyProg %d\n", dwErr);
}
else {
     if(WaitForSingleObject(pi.hProcess, INFINITE) == WAIT_FAILED) {
          dwErr=GetLastError();
          printf("Wait for object failed %d\n", dwErr);
     }
     dwValue=0;
     if(GetExitCodeProcess(pi.hProcess, &dwValue)){
          printf("Child returned %d\n", dwValue)
     }
     else {
          dwErr=GetLastError();
          printf("Unable to fetch return code %d\n", dwErr);
     }
}
```

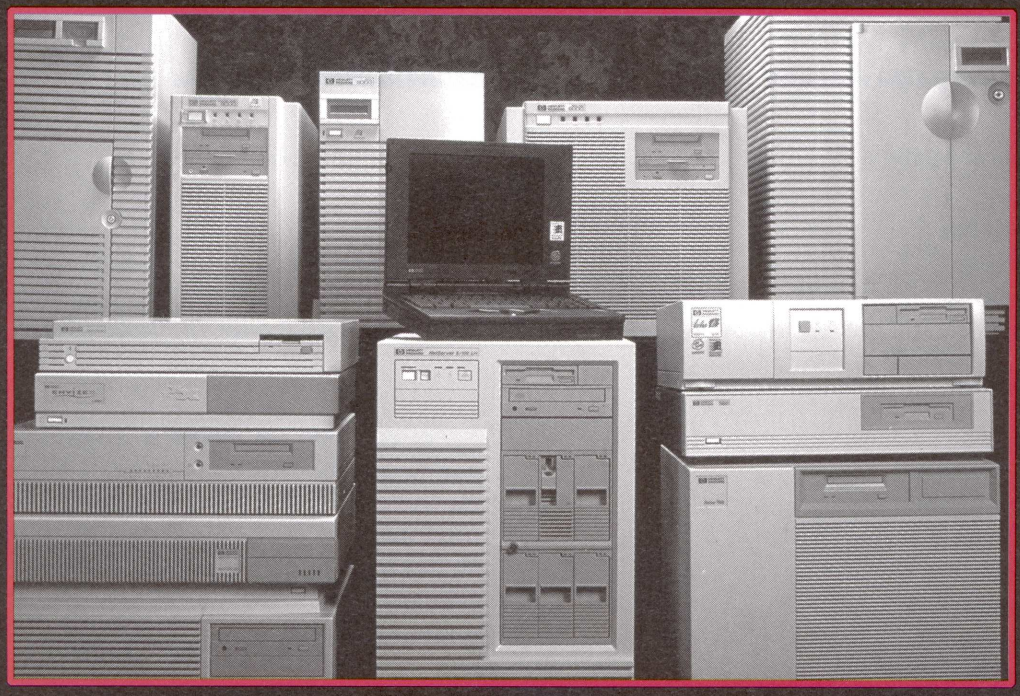
child process if passed the handle of the child process. We can get that handle from the process information structure returned to the parent; specifically `pi.hProcess`. The second parameter of the function is the time, in milliseconds, we should wait before timing out and declaring an error. I supplied the constant *INFINITE*, which specifies no time-out.

We pick up the returned value from the child by calling the *GetExitCodeProcess* after the child has finished. The first parameter is the child process handle we used before, and the second parameter is the returned value.

You might have noticed in *Listing 3* that `GetLastError` was called when the Win32 functions failed. Many of the functions only return a Boolean to indicate whether they succeeded or not. The function `GetLastError` can then be used to fetch the error code value.

The code I've just described will work for you, but I'd suggest you examine the `CreateProcess` documentation to find out about the other options available.

HP® SALES • RENTALS • SERVICE
1-800-422-4872



WORKSTATIONS

B Series
C Series
J Series
Visualize Graphics

COMPLETE LINE 700 & 800 SERIES IN STOCK!

BUSINESS SERVERS

E Series
D Series
G - H - I Series
K Series

AND MORE...

Add On Memory
Disc Arrays
Opticals
Netservers

Test & Measurement
Personal Computers
Data Acquisition
Notebooks

X Terminals
Printers
Plotters
Disc Drives

HEWLETT® PACKARD
Authorized Rental Company

VISA®
MasterCard
AMERICAN EXPRESS Card

THE RIGHT EQUIPMENT. RIGHT NOW.

TSA

2040 West Sam Houston Parkway N. • Houston, Texas 77043

713/935-1500 • Fax: 713/935-1555 • http://www.tsa.com • E-mail: info@tsa.com

CIRCLE 78 ON READER SERVICE CARD

LISTING 4 *Startup Information Structure*

```
typedef struct _STARTUPINFOA {
    DWORD   cb;
    LPSTR   lpReserved;
    LPSTR   lpDesktop;
    LPSTR   lpTitle;
    DWORD   dwX;
    DWORD   dwY;
    DWORD   dwXSize;
    DWORD   dwYSize;
    DWORD   dwXCountChars;
    DWORD   dwYCountChars;
    DWORD   dwFillAttribute;
    DWORD   dwFlags;
    WORD    wShowWindow;
    WORD    cbReserved2;
    LPBYTE  lpReserved2;
    HANDLE  hStdInput;
    HANDLE  hStdOutput;
    HANDLE  hStdError;
} STARTUPINFOA, *LPSTARTUPINFOA;
```

LISTING 5 *Process Information Structure*

```
typedef struct _PROCESS_INFORMATION {
    HANDLE hProcess;
    HANDLE hThread;
    DWORD dwProcessId;
    DWORD dwThreadId;
} PROCESS_INFORMATION, *PPROCESS_INFORMATION,
*LPPROCESS_INFORMATION;
```

Since I've mentioned threads, I'd like to show how easy it is for a process to create a thread. There is a Win32 API called `CreateThread` that looks somewhat like `CreateProcess`, but an even easier way to create threads is to use the C++ runtime library function `_beginthread`. Refer to *Listing 6*. The thread must be prototyped with the `UINT function (LPVOID)` ; syntax; also written *unsigned int function(void*)*; . The thread is then created by calling `_beginthread` with three parameters: the name of the thread function, the stack size, and a `void*` parameter. A stack size of zero means use the same stack size as the primary thread. Any type parameter may be passed to the thread, as long as it is typed within the routine, after passing the parameter as a `void*`. ◧

LISTING 6 *Thread Example*

```
UINT Count(LPVOID);
void main()
{
      int iValue;
      . . .
      _beginthread(Count, 0, (void*)iValue);
      . . .
}
// Worker thread (as opposed to a user-interface thread)
UINT Count(void* iVal)
{
      int I;
      I=(int)iVal;
      . . .
}
```

Bob Combs is the Director of Systems Architecture at PCSI® in Englewood, New Jersey, a company specializing in client-server technology. He is a Microsoft Certified Systems Engineer (MCSE) and holds a master's degree in electrical engineering.

MINISOFT IS

TOTAL CONNECTIVITY

MS

HP700/98, VT320, TN3270 & TN5250

NS/VT, TELNET & LAT

TCP/IP INCLUDED

ONLY $169

COMPETITIVE UPGRADES JUST $79

WINSOCK COMPLIANT • KEYBOARD REMAPPING • 32 BIT ARCHITECTURE
• DDE & OLE AUTOMATION • NETWORK FILE TRANSFER (FTP & NFT)
• RUNS ON DOS, MACINTOSH, WINDOWS, WINDOWS 95 & NT

800/682-0200

Internet address: sales@minisoft.com

MINISOFT

312 Maple Ave. Snohomish WA 98290
European sales: MiniSoft Marketing AG 41.41.340 23 20

CIRCLE 96 ON READER SERVICE CARD

MiniSoft 92 is a registered trademark of MiniSoft Inc. All other trademarks are the property of their respective owners.

Internet Goodies

by Joseph Berry

In a previous issue of *hp-ux/usr*, I described a program called *graphmaker* by Fabrizio Pivari. Mr. Pivari has sent me a note updating the address for his software. His current site is *http://www.geocities.com/CapeCanaveral/Lab/3469/index.html.*

Also in a previous issue, I discussed Sun's Java Development Kit (JDK). At the time, the latest version was 1.0.2. The new release of Java went out the door about a month ago. It is version 1.1.

ALT.SOURCES.D

gnumach (v 1.0)

Just when you thought you had enough toys to play with, the Free Software Foundation releases yet one more. If you enjoy hacking a little bit and/or have had fun with Linux, then you might want to look at yet one more UNIX clone, *gnumach.* This software can be found at all your favorite gnu sites including the following two:

ftp://prep.ai.mit.edu/pub/gnu/gnumach-1.0.tar.gz.
ftp://gatekeeper.dec.com/pub/GNU/gnumach-1.0.tar.gz.

This distribution is only for i386, i486, i586 (pentium), and i686 processors on PC-AT compatible machines.

The authors are looking for volunteers interested in ports to other architectures.

This kernel supports most ethernet cards and disk controllers. Interface cards for which Linux drivers exist can be ported easily. Non-network devices for which BSD drivers exist can be ported with a little effort, but not much.

mtools (v 3.6)

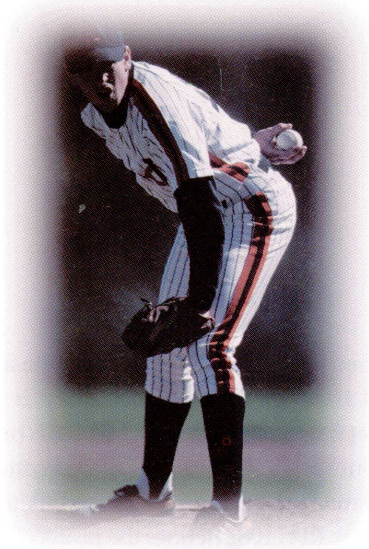
And again courtesy of the Free Software Foundation, a new version of *mtools* has just been released. *mtools* is a tool that supports floppy disk I/O on UNIX computers. In particular, it works well with PC floppies (including Windows 95 style long file names).

mtools works on AIX, HP-UX, and Sun operating systems. Anyone using floppy disks either to or from a PC should examine this program. It's easy to build and easy to use. *mtools* is also available from standard gnu repositories including *gatekeeper.dec.com.*

MISC

samba (v 1.9.16p11)

Ever since I purchased a PC running Windows 95 for my family, I longed to be able to conveniently and inexpensively interconnect this computer to my UNIX workstations at home. Of course, I had ftp on my PC as it was packaged with the TCP/IP software I had purchased for it. But what I really wanted was NFS. And that was expensive to purchase; no freeware version of PC/NFS exists. That was when I learned about *Samba. Samba* is a free SMB client and server for UNIX (and other) operating systems.



Integrated EC Solutions

for Windows NT, UNIX, Midrange, Mainframe, PC

Web Site www.stpaulsoftware.com
Email info@stpaulsoftware.com
Sales 800.998.4EDI (4334)

Don't Be Thrown a Curve Ball, Hit a Grand Slam with St. Paul Software's EDI ...

St. Paul Software is the leader in electronic commerce integration solutions.

We pioneered the first UNIX-based EDI product in the late 1980s and now offer a complete line for managing your electronic commerce environment. For Windows NT®, UNIX, PC, Mainframe *(MVS)* and Midrange *(HP3000, DEC VAX VMS and DEC ALPHA OPEN VMS)* environments. St. Paul Software's innovative EC/EDI software products, combined with our EC Center service bureau, professional services and value-added network options present you with a total solution.

As a leader in application integration, we can provide you with a seamless integration to Oracle, SAP, CA-MANMAN, SBT and other applications.

Hit a grand slam by integrating EDI with your applications!
Call us to find out why St. Paul Software should be your EDI partner.

Ask about our FREE CD-ROM!

ST. PAUL SOFTWARE
1450 Energy Park Dr, St. Paul, Mn 55108
Tel 612.603.4400 Fax 612.603.4403
"We'll Get You There™"

SAP R/3 COMPLEMENTARY SOFTWARE CERTIFIED INTERFACE
ORACLE Business Alliance Programme

CIRCLE 102 ON READER SERVICE CARD

FREE FOR HP USERS!

IMAP4/LDAP Enabled Email

Five free Motif, Windows and / or MAC clients.

Fax or mail completed coupon to receive download instructions and license keys. Valid until September 30, 1997.

Name
Title
Company
Phone
Fax
Email

SIMEON
Simplicity * Strength * Security

The ESYS Corporation
#900, 10040 - 104 St.
Edmonton, Alberta
Canada, T5J 0Z2

Phone: 403 424 4922
Fax: 403 424 4925
Email: sales@simeon.com
www.simeon.com

CIRCLE 47 ON READER SERVICE CARD

So what is SMB? Quoting from the README file for this package, "... it is the protocol by which a lot of PC-related machines share files and printers and other information such as lists of available files and printers." Using the software from the PC side is easy: You simply choose "network drive" and enter the UNIX host and path name. From that point on, your new drive (such as E:) points to the subdirectory on your UNIX computer.

On the UNIX side, you need two agents running, *smbd* and *nmbd*. Building them was easy. The required configuration file wasn't difficult either once I read the accompanying documentation. The software was developed by Andrew Tridgell in Australia. It is available via the Web at *http://samba.canberra.edu.au/pub/samba/* or via ftp at *samba.anu.edu.au* in the */pub/samba* directory.

sendmail (v 8.8.5)

The sendmail program is probably one of the most important UNIX utilities. Without it, the overwhelming majority of UNIX systems could not receive or deliver e-mail. Replacement programs do exist (I used to use *smail* on an AIX system), but most people stick with the original.

Unfortunately, there have been a number of security scares associated with the sendmail program, some valid and some not. Nevertheless, given its importance in the UNIX environment, you may not want to wait until your vendor supplies you with a more up-to-date version of sendmail with their operating system. Why not get the latest and greatest? Sendmail is available via ftp from *ftp.sendmail.org*. While you are there, you might also pick up a new copy of the Berkeley *db* package, which sendmail uses.

qpop (v 2.2)

For better or worse, I have recently established e-mail accounts for everyone in my family. That includes my four children aged 10 to 18. Their e-mail access is via a Windows 95 PC-based package but the Internet connection and mail server is on my UNIX computer. This is where I learned about the POP (or Post Office Protocol) server software. The client program (the PC e-mail software) uses POP to talk to the mail server and download pending e-mail messages. RFC 1082 defines the protocol used.

The latest and greatest version of POP (also called popper) is supported by Mark Erikson (*mark@qualcomm.com*). It was easy to build and straightforward to install. The software is available via anonymous ftp from *ftp.qualcomm.com* in the */quest/unix/servers/popper* directory and is called *qpop2.2.tar.Z*.

libdes (v 4.01)

A useful library to have in one's bag of tricks is *libdes*. This package includes a DES encryption library and DES encryption program. It supports ecb, cbc, ofb, cfb, triple ecb, triple cbc, triple ofb, triple cfb, desx, and MIT's pcbc encryption modes and also has a fast implementation of crypt(3). Typical uses include file encryption and programmatic encrypted transfer of data across the Internet.

The library is copyrighted by Eric Young (*eay@mincom.oz.au*) and is available from *ftp.psy.uq.oz.au* in the */pub/Crypto/DES* directory. It is also found on many other sites.

enscript (v 1.4)

Over the years, I have seen many programs that convert ASCII files to PostScript for printing on PostScript printers. Each seems to have its advantages and disadvantages and I have jumped from one program to the next. One day it occurred to me to see whether the Free Software Foundation (the GNU people) had a program to meet my needs.

I was surprised to find that I had never previously heard of *enscript*. The README file states, "GNU *enscript* is a drop-in replacement for the *enscript* program". Somehow I had missed this program (the original as well as the GNU replacement) all these years. *enscript* includes every option I have wanted plus more than I can ever use. My typical use of *enscript* is for dumping the contents of a file to my laser printer two-up, that is, two logical pages on one physical page. I define an environment variable, ENSCRIPT, as follows:

```
export ENSCRIPT="-fancy-header-borders-columns=2 --landscape-output=-"
```

To print a file, I simply type `enscript filename | lp` and the file is printed in landscape mode, two-up. You can get *enscript* from your favorite GNU repository site such as *gatekeeper.dec.com*.

COMP.LANG.JAVA.ANNOUNCE

Telnet Applet

Much valuable information can be learned from studying the code to this Java applet, a fully functional telnet application that includes terminal emulation support

for VT320/ANSI. The source code can be used to understand telnet negotiation, terminal emulation, and socket I/O. I learned more from this applet than from studying a number of Java books.

The software, including sources, is available via the web from *http://www.first.gmd.de/persons/leo/java/Telnet.*

COMP.LANG.PERL.MISC

zfilter (v 2.7)

I am now getting two or three pieces of electronic junk mail a day. It is starting to bother me. This program, *zfilter*, couldn't have come at a better time. As the README file says "... ZFilter is a vastly superior e-mail filtering program that can ... automatically destroy junk mail, chain letters... (it) can track and maintain logs of who writes to you the most, what actions it has taken the most, and up-to-the-second reports of how much mail has been received during the current year, month, day, hour, and minute."

The software requires Perl (which of course everyone has). It is available from the author's (Steven Zeck) home page at *http://www.guam.net/home/viper/files.html.*

COMP.UNIX.SOLARIS

xcoral (v 3.0)

Now this is a nice program! *xcoral* is a multiwindow, mouse-based text editor for editing various types of files for use in an X windows environment. Developed in France by Lionel Fournigault, Bruno Pages, and Dominique Leveque, this program is cognizant of C, C++, Java, Latex, HTML, RCS formats, and other formats. The program includes various display fonts, menus, a toolbar, a macro facility, and an undo mechanism. Included is a comprehensive online manual as well as a 175-page PostScript-ready manual. The editor will also color your text according to the type of document you are editing. Last but not least, *xcoral* includes *Smac*, a small, built-in ANSI C interpreter for dynamically extending the editor's functionality.

Building the software was easy. It has been tested on the following systems: SunOS, Solaris 2, Linux, AIX, HP-UX, IRIX, and OSF/1. The software is available from the contributors at *ftp.inria.fr* in directory *X/contrib-R5/clients* or *ftp.x.org* (*contrib/editors*). If you are looking for a nice graphical text editor for a UNIX system, spend some time with this program. It's worth the investment.

WEB PAGES

http://www.windows95.com

I apologize to all the UNIX-only purists for mentioning this site, but I find that more and more it is not possible to remain a UNIX-only purist. In my own office at Landmark, I have a Solaris X.86 workstation as well as a Windows 95 workstation. I use them both. At any rate, for the person needing utilities (including UNIX-lookalike software) for Windows 95, this is a great site. Lots of good shareware, too. ◧

Joseph Berry is a senior software developer at Landmark Systems Corporation in Vienna, Virginia. He is one of the authors of Landmark's PerformanceWorks products, PerformanceWorks/Smart Agents for UNIX. A former HP 3000 systems specialist for Hewlett-Packard, he has been in the computer industry for more than 25 years. He can be reached at joe@wayne.unix.landmark.com.

Workstation DIRECT
Upgrading The World

Operating Environment:
HP 9000 Series; B Class; C Class; D Class; J Class; K Class; Series 800 E, F, G, H, I; Series 700, HP 3000 9x9KS; Series 900 models 9x7, 9x7LX, 9x7RX, HP Vectras, HP NetServer, HP Printers: all configurations.

Product Description:
Workstation Direct, a pace-setter in the electronics industry, manufactures high-end memory products for workstations, servers, and PC's. Our products are consistently reliable, meet or exceed OEM specifications, carry a lifetime warranty, and are 100 percent compatible within most computer environments. We offer over 125 memory expasion modules including HP, Sun, SGI, IBM RISC 6000, and popular module configurations for workstations, servers, PC's and printers. Our modules are built to the highest standards on multilayer boards and are assembled using premium grade parts. For high quality, low-cost memory, free technical support, and fast delivery, we are your source for memory.

Our distinctive blue printed circuit boards make it easy to identify Workstation Direct memory. All memory modules are assembled with IC's from major DRAM manufacturers. All finished products are 100% tested prior to shipment and a simple part numbering system makes for easy ordering. We offer next-day delivery across the United States with same day messenger service available in the metro NY area.

Workstation Direct
55 HollyHill Lane
Greenwich, Connecticut 06830
Phone: 1-888-620-1600
Fax: 1-800-371-2657
http://www.wksdirect.com

CIRCLE 121 ON READER SERVICE CARD

X-Watch

by Larry Headlund

Lisa with an S

In my last column I considered the design of the user interface of the Apple Lisa, the interface that evolved to the Apple Mac with all its influence. What was exciting for the developers was that they were presented with a blank slate. There was no installed base for graphical environments, no user habits or hard-won user knowledge or loyalties to conform to. There was not even an existing paradigm to guide or restrict the developers. So successful was the fruit of their efforts that no one has had such an opportunity since.

Such an influential design creates in me an interest in the design and testing methods used to create it. The article in *interactions* referenced below and the perspective of the last decade and a half give us a chance to critique both the result, the Lisa/Mac interface, and the process of producing it. Since X was having its beginnings in the early 1980s overlapping the 1978 to 1983 development period of the Lisa, I can contrast the decisions made in each case.

Lisa Goals

Like all good development projects, the Lisa started with a set of goals and guiding principles. In the Lisa's case they were:

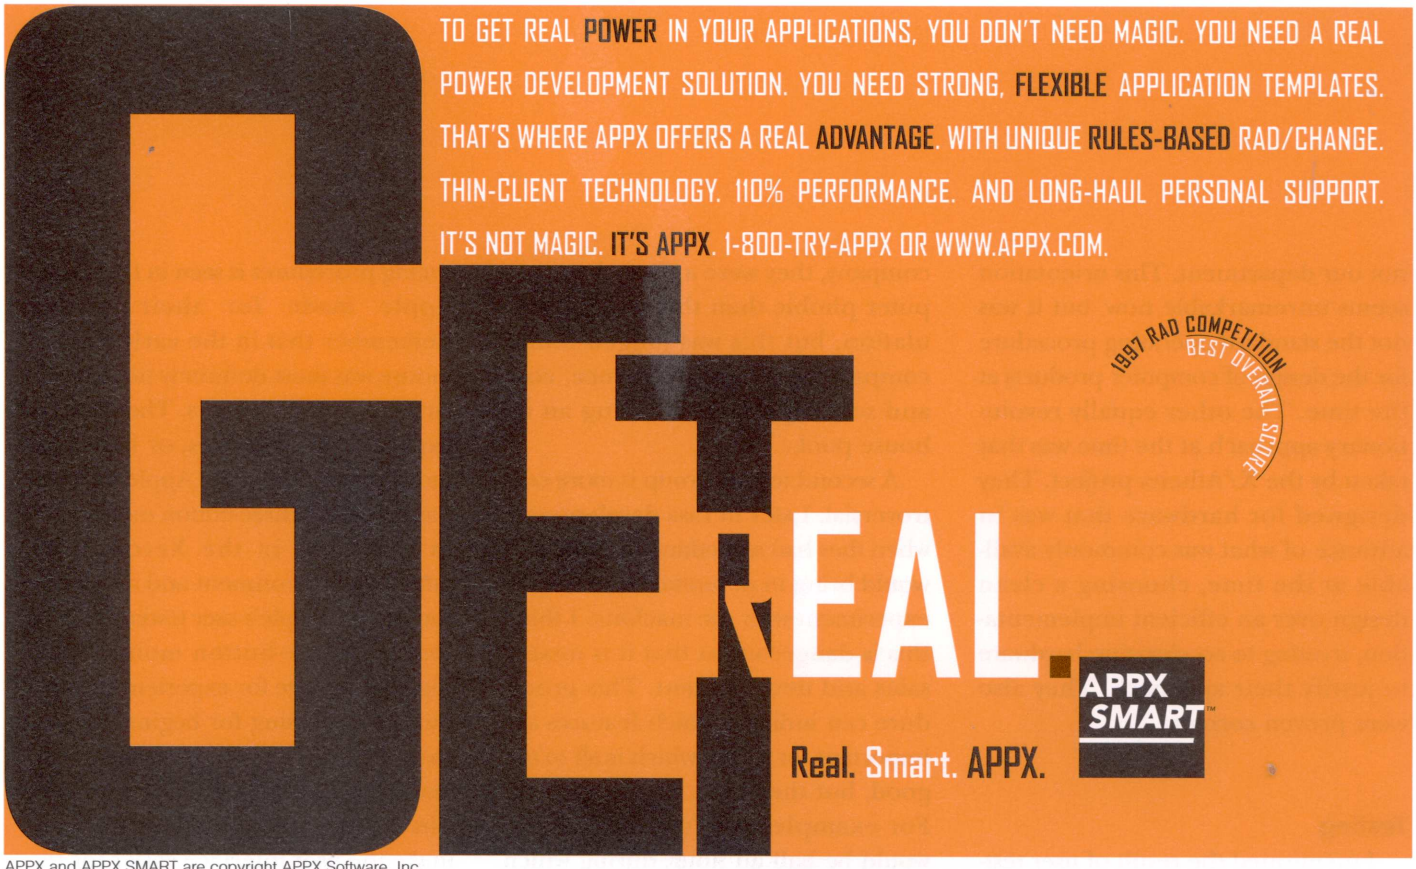
1. Lisa must be fun to use.
2. Lisa will require extremely minimal user training and "hand holding."
3. Lisa will adhere to the concept of gradual learning.
4. Lisa will use graphics aggressively.

There were other goals in line with Apple's intention that the Lisa be a general-purpose business machine that would allow user customization and enable the user to handle several tasks at once, with consistent user interfaces and consistent error and warning messages. It is the four goals enumerated above that I want to home in on.

The first and fourth goals, fun and aggressive graphics, were marketing decisions to differentiate the Lisa from the competition. They meant that when a choice had to be made between equal or nearly equal approaches, the nod would go to the approach with the most flash. An example of this kind of decision was the design of the filing interface.

One filing interface considered the design team called 'Twenty Questions." In this model, when a user wished to access a document he would be prompted to choose a document from a list. He would then choose an action to apply to the selected document(s). Another competing interface was the now familiar one with icons and full-size documents and dragging and dropping for actions. The two approaches were a dead heat in user testing in such areas as speed of learning, speed of task completion, level of comprehension, and error rate. However, no user preferred the "Twenty Questions" design and some preferred the iconic approach because it was "more interesting and fun." Hence the nod went to the iconic design.

Note what was not included in the decision: nothing about computational efficiency. If the hardware could do it, and the Lisa team was taking advantage of the then new Motorola 68000, that was sufficient. Nothing about ease of programming, that was



APPX and APPX SMART are copyright APPX Software, Inc.

CIRCLE 35 ON READER SERVICE CARD

Consider the advantages of the INTEREX DIRECT MAIL SERVICE

- ✦ 46,000 Hewlett-Packard computer users in our database
- ✦ A one-week turnaround time
- ✦ Less than 2% returns because of incorrect addresses
- ✦ 74% of our users are their company's decision makers
- ✦ Full refund for all nixes (bad addresses)
- ✦ 15% discount on all orders over 10,000 labels
- ✦ List updated weekly

We offer the most value for your money through direct mailings and the INTEREX Card Deck.

Contact Phil Nguyen, INTEREX,
P.O. Box 3439, Sunnyvale,
CA 94088-3439, USA

(800) INTEREX • (408) 743.4619 FAX; (408) 747.0947

not our department. This orientation seems unremarkable now, but it was not the standard operating procedure for the design of computer products at the time. The other equally revolutionary approach at the time was that taken by the X/Athena project. They designed for hardware that was in advance of what was commonly available at the time, choosing a clean design over an efficient implementation, trusting to accelerating hardware to justify their approach. They also were proven correct.

Testing

I mentioned the result of user testing above. The structure of user testing is at least as important as goals. Your testing procedures must align with your goals if your goals are to be measured. It does no good to measure weight if your goal is speed. In the case of the Lisa, a key way to measure progress toward the middle two goals, minimal user training and gradual learning, was very abbreviated testing with potential clients and typical users. What the Lisa team meant by these two goals was that a new user should be able to accomplish simple tasks immediately with no instruction and that knowledge of how to do more complicated tasks would be acquired gradually on an as needed basis. I doubt that anyone was ever told to read the fine manual.

The first pool of test subjects were new Apple employees with no computer experience. There was the reasonable assumption that these employees could represent the pool of potential users. Perhaps, since they were coming to work for a computer company, they were a shade less computer phobic than the general population, but this was judged minor compared to the convenience, cost, and security factors of using an in-house pool.

A second testing group is more controversial. Later in Lisa development, when they had something to show, they would bring in potential customers to experiment with the machine. I think this is dangerous in that it is mixing sales and development. This procedure can indicate which features are important for a sale, which is all to the good, but the result may be skewed. For example, the "testing" period would be half an hour, during which time the customers could see if they could perform simple tasks without help. This fits in nicely with the paradigm of gradual learning. However, are sales in the real world based on 30-minute demonstrations? If the answer is yes for this product, then you have a powerful tool in this testing for telling you what a successful product should look like. If the answer is no, then you may be hurting yourself by thinking that a successful demonstration brings you a sale. This is because the thirty-minute demonstration is so different from the actual function of the machine. A computer interface is what a purchaser will interact with for a good portion of his or her working life. Is the impression after the first 30 minutes the best measure? I am reminded of the consumer interfaces that "any idiot can use," and after five minutes only an idiot would want to.

The One-Button Mouse

The result of these design goals and testing procedures is seen in the choice Apple made for their mouse. Remember that in the early 1980s a mouse was most definitely not a familiar object to consumers. They had no expectations, practices, or habits for mouse manipulations. Apple's choice was between a three-button mouse such as that used in the Xerox PARC Smalltalk environment and a one-button mouse. Apple's user testing showed that the three-button mouse had a slight advantage for experienced users but was confusing for beginners. The extra buttons also made it harder to learn the Lisa user interface quickly. The one-button mouse was chosen to make the user interface easier for the first-time user.

This is a real designers' decision. You are a first-time user once and a gradually more experienced user for the rest of your life. However, you don't want to spend too much time learning before you can get your work done. The decision of whose interest to favor, neophyte or master, is a judgment with no right answer. Note how the design goals and the testing procedure, particularly the thirty-minute tests with potential customers, all bias the decision towards the neophyte. Famously, the X universe went with the three-button mouse.

What You See Is All You Get

We go now from the detail of the mouse design to the grandness of the metaphor. To use effectively and in some sense understand an interface, the user must have a metaphor, a mental picture of what is going on. It helps if the metaphor bears some relationship to what is really going on under

the hood, to avoid jarring inconsistencies and to enable the user to make accurate predictions of the consequences of an action. The Apple Lisa developers chose the now familiar desktop metaphor along with analogies to file cabinets and the famous waste basket. This has the advantage of familiarity to the user, hence allowing gradual learning. However, any metaphor is also limiting. The computing environment takes on some of the limitations of the actual desktop. Documents can be misfiled, lost under other documents on the desktop, and so on.

The computer desktop can do some things a physical one cannot, for example automatically keep a list of all documents on the desktop and bring any one of them to the forefront at a button click, but the metaphor still gets in the way. This is what is meant when we say that a GUI makes simple tasks easy and complex ones impossible. Think of the tasks that are simple with shell pattern matching or with the UNIX *grep* command that are tedious with a GUI. Since this desktop metaphor permeated the design of the Lisa, the developers had a difficult time coming up with a functional analog of the UNIX *find* command, for example. This metaphor spread beyond the Lisa and Mac so that it was almost a decade later that the World Wide Web made popular some of the modes of operation that only a computer could support, modes and functions that were outside the original paradigm.

Conclusion

History, even in so new a field as computers, is always worth studying. It can show us both how and why we are where we are. ◧

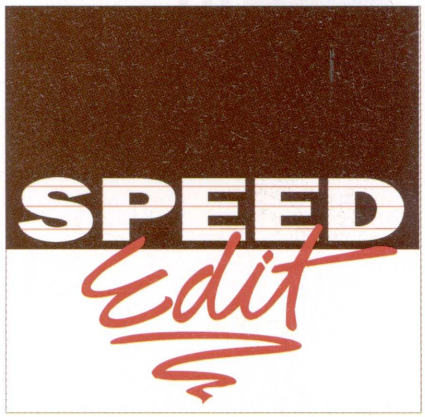
Larry Headlund is the president of Mathematical Engineering, Inc., a UNIX and Motif development company. He has been working with commercial UNIX since 1983 and with X since 1987. He can be reached at lmh@world.std.com or at 1.617.242.7741.

References:

"Inventing the Lisa User Interface," Roderick Perkins, Frank Ludolph and Dan Smith Keller, *interactions*, January-February 1997, Volume IV.I

"What you see is ALL you get," Harvey Lehtman, *interactions*, January-February 1997, Volume IV.I

Client/Server Text Editing



PC Editing for MPE & UNIX

Now MPE and UNIX users can transparently edit their host based files right on their PC with a fast, easy and efficient Windows text editor, in much less time than traditional host based text editors. Best of all, you get all the powerful features found in the popular SpeedEdit text editing subsystem, while at the same time offloading the entire editing task from your host system. It's like upgrading your host computer for a fraction of the cost, and getting the worlds best text editor as a bonus.

SpeedEdit is loaded with time saving features, over 200 of them. SpeedEdit is language sensitive, completely configurable and works in both networked and non-networked environments. Native versions of SpeedEdit are available for MPE, MPE/iX, and most popular UNIX systems as well as the Windows based Client/Server version. Client/Server works with just about any host computer and even allows you to edit files over the Internet.

Visit us on the WEB, call, write or email for a **free demo**.

Inclination Software, Inc.
PO Box 8668
Incline Village, NV 89452
(702)831-5595 Fax (702)831-4979
http://www.isiinc.com
email: sales@isiinc.com

CIRCLE 51 ON READER SERVICE CARD

HP-RT

by Anita Harris

This month's column will address questions that might arise when writing an HP-RT driver.

Q: I am writing an I/O driver that will be writing data to an HP-RT swsm on my local processor. I have the swsm defined in my *CONFIG.TBL* as:

```
0:local_mem:A32+DATA_ACCESS:0x20001000:0x1000:2:SYSTEM_ONLY
```

My CPU number is 2 and I have 64 MB of local RAM. My driver and kernel make without errors but when I boot my kernel, I get the message:

```
** Failed to create an inode for user-level SWSM local_mem
```

If I take the swsm out of my *CONFIG.TBL* and remake the kernel, I don't see the error message. What is wrong with my swsm declaration?

A: HP-RT uses direct mapping when slave mapping A32 address space, as you are doing by defining a swsm on your local processor. This means that it has 256 MB of VME A32 address space mapped directly to its lowest 256 MB of memory. The first byte of the VME region maps to the first byte of physical memory on the HP-RT system. The kernel will utilize the lowest physical memory addresses. This area is from PA address 0x00000000 to approximately 0x00900000. These addresses then have VME addresses of 0x20000000 to approximately 0x20900000. You can define your swsm at a VME address above 0x20900000 to avoid a conflict with kernel space.

Q: I am writing and testing a driver using two different variations of interrupt vectors. One variation uses second-level interrupts. Another uses a first-level interrupt scheme. I find that when I boot a kernel made using the first level interrupt scheme, I get errors indicating that I am using second-level interrupts. I checked my driver make files and the only difference I see is in the section that locates the *drvrinfo.o* file. The sections in question are below. Is this causing my problem?

```
CCRT    =  $(HPRTroot)/hpux/bin/ccrt
DRIVER_SRC      = /users/me/mydriver
TRG_DRIVERS     = $(HPRTroot)/etc/conf/drivers
TRG_DEVICES     = $(HPRTroot)/etc/conf/info
TRG_DHEADERS    = $(HPRTroot)/etc/conf/dheaders
```

Second level interrupt makefile:

```
$(TRG_DEVICES)/skelinfo.o : $(DRIVER_SRC)/skelinfo.c \
                            $(DVR_HEADERS)
        $(CCRT) $(TRG_DEVICES)/skelinfo.c \
                -I$(DRIVER_SRC) -I$(DRIVER_SRC) \
                $(CFLAGS) -o $(TRG_DRIVERS)/skelinfo.o >skelinfo.lst
```

First level interrupt makefile:

```
$(TRG_DEVICES)/skelinfo.o : $(DRIVER_SRC)/skelinfo.c \
                            $(DVR_HEADERS)
        $(CCRT) $(TRG_DEVICES)/skelinfo.c \
        -I$(DRIVER_SRC) -I$(DRIVER_SRC) \
        $(CFLAGS) -o $(TRG_DEVICES)/skelinfo.o >skelinfo.lst
```

A: The section of your makefile for the second-level interrupt places the *drvrinfo.o* file in the *$(TRG_DRIVERS)* directory. The effect is seen when a different driver with the same *drvrinfo.o* name using different interrupt vector information locates the *drvrinfo.o* file in the recommended *$(TRG_DEVICES)* directory and the other file remains in the *$(TRG_DRIVERS)* directory. The kernel uses the *drvrinfo.o* file to bind key driver information at system installation. Placing the *drvrinfo.o* file in the *$(TRG_DRIVERS)* directory forces the entry points to be resolved from this file first rather than from the expected location in *$(TRG_DEVICES)*. So, when you made your kernel using the first-level interrupt scheme, the *drvrinfo.o* information used to resolve entry point was from the second-level interrupt *$(TRG_DRIVERS)/drvrinfo.o* file.

Q: I have written my HP-RT driver and I am now configuring it into the kernel by inserting

```
I:mydriver.cfg
```

at the beginning of my *CONFIG.TBL* file. I had a problem with my driver during boot with a swsm configuration so I removed the line from the *CONFIG.TBL* file and remade the kernel to make sure I could still make and boot a kernel. The same swsm problem remained! I confirmed that I had removed the entry from the *CONFIG.TBL* file. Now, nothing I do removes this problem from the new kernels that I make. What is wrong?

A: This problem may be the result of inserting your driver at the beginning of your *CONFIG.TBL* file. New drivers should always be placed at the end of the file. Refer to the *Driver Writing in the HP-RT Environment* manual, Chapter 5, "Installing Drivers," Figure 5-1, "Typical Driver Configuration Flow Diagram." You can see that the information provided in the *CONFIG.TBL* file is used by the *config(1)* utility to produce the nodetable *nodetab*. It contains the names and descriptions of the device files in your kernel. Each major I/O device on a system is numbered, by the system, according to its position in *CONFIG.TBL*. Placing a new device at the end of the *CONFIG.TBL* file avoids changing major numbers for existing drivers. The *CONFIG.TBL* file is also used to create the *CONFIG.h* include file, which is used to configure the driver and swsms into the kernel.

If these files are corrupted by the insertion of your driver at the beginning of the *CONFIG.TBL*, file you will need to purge the *nodetab* and *CONFIG.h* files and remake your kernel to resolve the corruption.

Q: I am attempting to debug my driver using *pdb*. I have an RS-232 connection from */dev/tty0p0* of my HP-UX host to port B of my HP-RT target system. I boot my 743rt kernel using:

```
ISL> rtboot -aK -rramdisk /mykernel
```

to debug the install entry point of my driver at bootup. I run *pdb* from my 10.20 HP-UX host using:

```
$ $HPRTroot/hpux/bin/pdb -B -f/dev/tty0p0
                                /home/tftpdir/mykernel
```

I receive the error message:

```
Unable to open communication link to target system
/HP-RT/hpux/bin/realpdb: Permission denied
      open(/dev/tty0p0,2) failed
```

If I make myself root, I can successfully run *pdb*. What do I need to do to run *pdb* without being root?

A: You can see from the error message that the open of */dev/tty0p0* failed. In the *HP-RT Symbolic Debugger User's Guide*, Chapter 3, there is a section titled "Setting Up the HP-UX Terminal Device File." It states that the access permissions on the device file must be set to read and write, that is *rw-rw-rw-*. If write permission is not set for *group* and *other*, only

the *owner* of */dev/tty0p0*, root, will be able to run *pdb*. As root perform:

```
# chmod 666 /dev/tty0p0
```

and you should be able to run *pdb* as a non-superuser.

Q: I am attempting to debug my driver using *pdb*. I am getting the message:

```
Passed through init_rs232_card().
/HP-RT/hpux/bin/realpdb Terminal doesn't support window mode
```

I am using a CDE dtterm window on my HP-UX 10.20 host. What do I need to do to get *pdb* to work?

A: Again, you can refer to the "Setting Up the HP-UX Terminal Device File" section of Chapter 3 in the *HP-RT Symbolic Debugger User's Guide* for assistance. If you wish to use a terminal emulator and you can use an hpterm window instead of the dtterm window, you will be more successful running *pdb*.

Q: I have a 742rt HP-RT processor running revision 2.21 of HP-RT. I want to transfer data to my I/O card using DMA. I understand that there is no DMA engine on the 742rt, but my I/O card can master A32+DATA_ACCESS DMA transfers. What do I need to do in my local process to set up the transfer?

A: You have two options. You can describe local memory on your 742rt processor in a swsm area. The beginning VME address of this contiguous memory and its size can then be passed to the remote card or processor. This is the easiest way.

If you do not have available memory or cannot use it for defining a swsm for DMA, you can translate virtual addresses from your driver to physical addresses using *mmchain(9)* or *mmchainjob(9)* functions. The starting physical addresses and sizes of all non-contiguous memory segments that make up a contiguous virtual memory segment are written into an array of type struct *dmachain*.

```
struct dmachain {
      long address;
      int count;
      };
```

The array looks something like this:

	Physical Address	Size
dma_array[1]	0x10000	0x1000
dma_array[2]	0x13000	0x2000
dma_array[3]	0x16000	0x1000
	...	...

This effectively chains memory segments together to make them contiguous. This array can then be used to set up DMA by the remote device. The remote device will expect a VME address to read from, and a count or size of the buffer to read. You will need to utilize the fact that A32 space uses direct slave mapping for swsms. You should define a "dummy" swsm in A32 space. This assigns an A32 region to your CPU, which then allows you to calculate a VME address for each physical memory address in your system.

For example, let's say we are developing a write entry point for our driver. It will look something like:

```
mydriverwrite(myDrvrData, fd, buffer, buffcount)
mydrvr_data *myDrvrData;
struct file *fd;
char *buffer;
int buffcount;
```

We set up our array for doing the dma chaining:

```
struct dmachain dma_array[100];
int num_elements;
```

In our code we call *mmchain* to create the dma_array chain of physical addresses for our use.

```
num_elements = mmchain (dma_array, buffer, (long)buffcount);
```

We will do a DMA transfer for each element in the dma_array chain for the size of the segment available. So we pass the address of each element calculated as a VME address to the I/O card.

Continued on Page 86

HEWLETT-PACKARD AND LOTUS DOMINO

BUSINESS SOLUTIONS TO INTEGRATE THE WEB



EXTENDING YOUR BUSINESS

Stop by the Lotus® Domino™ demo in the Hewlett-Packard® booth and enter to win: a Domino Snowboard, a copy of Domino or a copy of Weblicator™!

Your organization knows what it wants to do. Tap the power of the Internet for distributing up-to-the-minute information. Use your intranet as the foundation for real collaboration throughout the enterprise. Foster closer communication with your customers and partners via your Web site. No matter what your vision, Hewlett-Packard and Lotus can take you there with powerful, Internet-ready HP servers and the unequaled capabilities of Lotus Domino — the world's most powerful, interactive Web server technology. For powerful intranet/Internet solutions, the choice is clear — Lotus Domino on Hewlett-Packard servers.

For more information visit our Web address @ http://www.hp.com or http://www.lotus.com

HEWLETT® PACKARD

Lotus .domino

CIRCLE 45 ON READER SERVICE CARD

Lotus is a registered trademark and Domino and Weblicator are trademarks of Lotus Development Corporation.
© Copyright Hewlett-Packard Company 1997. All rights reserved.

```
for (i=0;i<num_elements;i++)
{
beginaddress = dma_array[i].address + A32_baseaddress;
size = dma_array[i].count;
...
/* Begin DMA transfer */
}
```

Q: How can I dynamically determine the base address of my VME A32 space region? I know that it is in region 2 because that is where I defined my A32 swsm, but I want my driver to be able to determine the value.

A: Here is a routine you can use to determine your A32 base address.

```
char *get_A32_base_address()
{
struct sys_shmem *swsm_pointer;
int i;
struct lowmem *lm;
extern struct sys_shmem *swsm_ptrs[]; /* Array of pointers to SWSMs
                                         defined for this VME system
                                         in CONFIG.TBL */
lm = LOWMEMBASE;  /* See /usr/include/machine/frame.h for the lowmem
                                         structure. */
/* Get the CPU number of the local processor from the lowmem
  structure. */
cpuNumber = lm->cpu_number;

/* For each swsm defined in this VME system... */
i = 0;
swsm_pointer = swsm_ptrs[i++];

while (swsm_pointer != 0)
    {
    /* Is this swsm defined on our local CPU? */
    if (swsm_pointer->cpu_number == cpuNumber)
       {
       /* Determine if it is an A32 swsm. See sysshmem.h
        for SPACE_ID macro. */
       if (SPACE_ID(swsm_pointer->space_access) == A32)

       /* VME address of beginning of swsm. Remember there
       can only be one A32 swsm per processor. */
       return (char *)(swsm_pointer->BUS_address & 0xF0000000);
       }

    swsm_pointer = swsm_ptrs[i++];
    }
return (0);
}
```

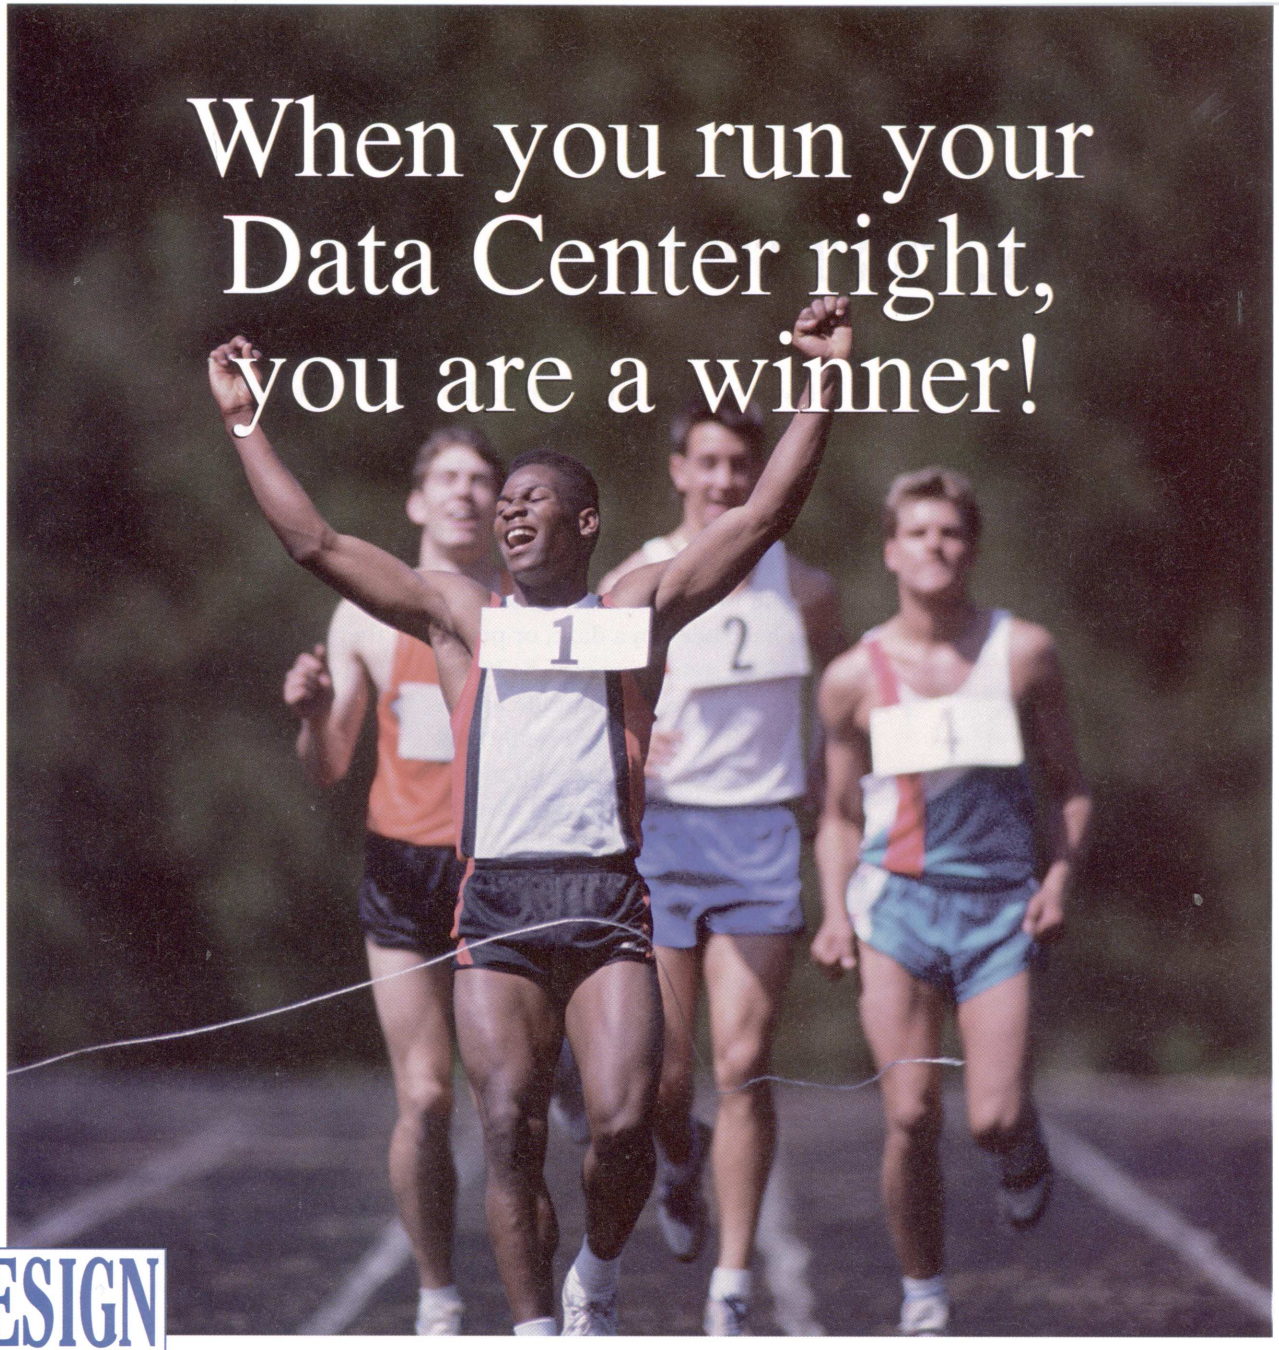
Important Patch Notice

There is a known problem documented in SR 5003361089 regarding timers using sub-10ms 'times' stalling timer expirations for up to 40 seconds. It is recommended that all customers install the appropriate patch for their OS revisions to avoid encountering this problem. The patch numbers for each currently supported revision are listed below. Contact your local HP Response Center to obtain the patch for your OS.

HP-RT OS Revision	Patch Number
2.12	PT970321_001
2.21	PT970325_000
3.0 + warranty patch	PT970401_000



HP-RT Operating System questions are answered by Anita Harris, a support engineer in the HP-RT Worldwide Technology Expert Center. She has worked with Real Time systems for eleven years as a customer, a Real Time Response Center Engineer, and HP-RT WTEC Engineer.



DESIGN 3000 Plus INCORPORATED

When your Data Center has problems it's a race against time until they are corrected. We have JMS to schedule and control your batch jobs, and CALLBACK to monitor and notify you of critical conditions. You will be notified by phone, pager or alpha pager of any user defined conditions or other problems that occur within your Data Center.

JMS Job Management System

JMS puts you in control of your Data Center with scheduling and execution of batch jobs. JMS has no command language to learn so it's easy to use and yet has sophisticated scheduling capabilities that are second to none. JMS network capabilities allow scheduling and job dependencies across the network. Restriction features allow you to limit those pesky system hogs. Ad hoc streamed jobs may be incorporated to allow complete batch job control.

Callback System Down Detector And Job Notification System

Voice/Beeper/Digital or our new Alpha Pager notification for: • Adverse Temperature • Power Loss • System Hangs • Job Aborts • Printer Status • Physical Conditions • User Requests • Logon Security • Reply Pending • Console Messages • Special Job Events • Spoolfile Scanning • And much more...

We are the leaders in tools for the "Lights Out Environment." Call us today for a free 30 day demo. MPE/iX or HP-UX.

Corporate Headquarters
3470 Pipebend Pl. NE • Suite 120 • P.O. Box 13086 • Salem, OR 97309-1086
Phone: (503) 585-0512 • FAX: (503) 585-1706 • E mail: Design3000@aol.com

International Sales Office
System Software Intl. • Oakmoore Court • Kingswood Road • Hampton Lovett
Droitwich, WR9 0QH, UK • Phone: +44 (0) 1905 794646 • FAX: +44 (0) 1905 794505

HEWLETT PACKARD Channel Partner

CIRCLE 44 ON READER SERVICE CARD



CSL Perspective

HAVE YOU EVER FELT really embarrassed ? Like you just wanted to bang your head against a wall and mutter to yourself, "Boy, that was a pretty stupid thing to do!" That's a pretty accurate description of the way I felt after reviewing my column from last issue. I must have been in a daze, or preoccupied with something else. So let me offer my sincere and humble apologies to each of you since I really blew it. Now let me try to make it up to you.

I've mentioned, for several months, the FAST project. FAST stands for Freely Available Software Technology. It is a collection of "freely available" software from various sources such as the Internet and HP. This collection will include the latest versions of many of the most popular software packages ported to HP-UX 10.10 and distributed on CD-ROM and through archive sites such as InterWorks and the HPPD Centre. The packages will be bundled together in SD format to allow easy installation and will be integrated into the Common Desktop Environment (CDE) as appropriate. Each package will include additional documentation, produced by the FAST team, which will give users and administrators additional assistance in effectively utilizing the software at their sites.

This is a major undertaking, to be sure, but the benefits to be gained in the long term are significant. For years one of the most frequently asked questions on the *comp.sys.hp.hpux* newsgroup has been, Where can I get an HP-UX version of _____ ? This question has been driven by users and system administrators alike, many of whom are using HP-UX for the first time.

More and more HP-UX systems are finding their way into sites that had no UNIX systems or are primarily non-HP UNIX systems. And with the major upgrade from version 9.X to 10.X, it has become very apparent to many long-time HP-UX users that there is a significant need for this software. This huge, unmet need we believe is an enormous opportunity not only for the various users groups but for HP as well. That is why HP is assisting us in this effort.

The FAST project team has established a process and some standard specifications for doing this porting work. Since the 10.10 version of HP-UX now uses the SVR4 file system layout, the standard strongly encourages the porters to utilize this structure. Applications will be built to install in the */opt/<appname>* directory, which provides a more manageable location for system admins. Each package will be compiled with HP's C compiler in ANSI mode wherever possible. Porters can also attempt a build using GNU C, but it is not required. The resulting binary packages may not use transition links and will be packaged in Software Distributor format for ease of installation and network distribution.

Because of the huge volume of software that needs to be ported, the team has come up with categories or bundles. Each category contains the most popular packages so that the first release covers a large portion of what is needed. The current list includes:

- Software Development
- Mail/News
- Networking Tools
- System Monitoring
- Word Processing

- Image Handling
- File Handling
- General-purpose Tools
- Alternate Un*x Tools
- System Administration
- Games

One of the value-enhancing aspects of this porting activity involves the documentation process. Many of us have found that the existing documentation is not very useful for real people. This lack of good documentation can be a significant hurdle to users and sys admins alike and may ultimately result in the user rejecting a useful package. To overcome this, the project team has hired technical writers to generate user-level documentation. The writers will interact with the porters in an interview format to get the basic factual information and then generate base-level docs. There will be some review and then final versions will be published in various forms, including HTML.

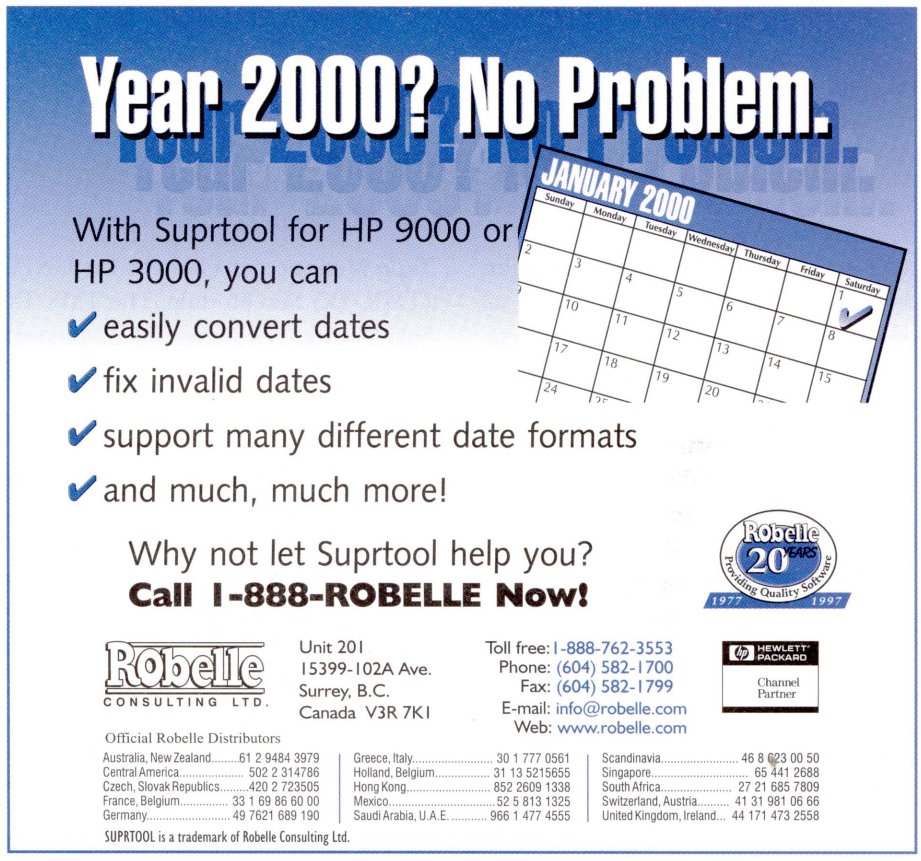
As for when you might be able to get ahold of the fruits of this effort, our current schedule calls for availability on July 1, 1997. This CD will be added to the HP Price List (Thanks to HP!) and the entire CD will be placed on the InterWorks, Interex, and HP-UX Porting Archive sites for download. We will also be looking at some special distributions for Interex On-Line (CSL) users during the year.

As a member of the team who has been involved in porting software for years, I'm extremely enthusiastic about this effort and want to acknowledge the extraordinary effort that Greg Cagle, Dave Shaw, and Doug Siebert have given to get this project off and running. Please drop them a note of thanks and also send one to the other volunteer members of the team:

- Ralph Dickman, Interex staff
- Debbie Lawson, InterWorks/Interex staff
- Greg Cagle, HP
- Jim Rice, HP
- Doug Siebert, University of Iowa
- Dave Shaw, InterWorks Librarian
- Dave Eaton, Honeywell
- Pierre Mathieu, Geomatics Canada
- Heath Kehoe, Norand
- Paul Gerwitz, HP and Interex CSL Chairman
- Chip Richards, Honeywell
- Jeff Hildebrand

Before I sign off, I'll mention that I've changed employers. I now work for Hewlett-Packard in the PSO, doing very much the same type of work I did at Kodak. You will find my new phone number and e-mail address below.

Paul Gerwitz is chairman of the Contributed Software Library committee and a technical consultant at Hewlett-Packard. For 16 years, he was a technology specialist and analyst at Eastman Kodak in Rochester, New York. He can be reached at 610-408-6526 or via e-mail at gerwitz@interex.org.



CIRCLE 116 ON READER SERVICE CARD

Continued From Page 15

OSF-certified Motif + Development, an industry-standard Linux distribution, all major European languages, and Japanese.

Users can customize their desktops without sacrificing administrative control over system configuration. It provides a common method of installing and integrating applications in a distributed multiuser environment.

CDE Business Desktop is priced at $299.

Contact Xi Graphics, phone: (303) 298-7478, fax: (303) 298-1406, e-mail: info@xig.com, http://www.xig.com.

Storage Management

P^2 Software, Inc. has announced its Tape Management System for UNIX (TMX/IX) Silo Module. The TMS/IX Solo Module uses StorageTek Automated Cartridge System Library Software and supports various UNIX platforms.

ACSLS provides connectivity to StorageTek Nearline tape libraries, including PowderHorn 9310, WolfCreek 9360, and TimberWolf 9710.

The TMS/IX Silo Module provides tape management and interface control for UNIX platforms attached to StorageTek Robotic Libraries. It gives the same functionality, control, and managed environment for UNIX systems in a data center as exists for the mainframe counterparts.

TMS/IX can manage any number of UNIX hosts connected to tape devices in the silo. The product supports and provides extensive validation for IBM or ANSI labels. Additionally, tape catalogue records from mainframe tape management can be imported into the TMS/IX tape catalogue, for permanent or temporary control.

Contact P^2 Software, phone: (214) 340-1833, fax: (214) 349-3765, e-mail: jstafford@p2soft.com, http://www.p2soft.com.

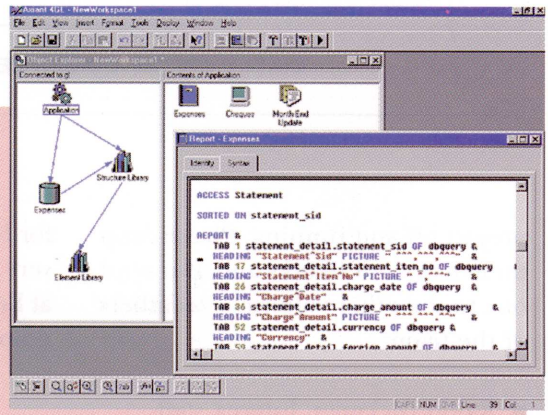
New from Onion Peel Software

RoboMap V2

Onion Peel Software has announced a complete redesign of RoboMap for OpenView. Integration with new OpenView 4.1 APIs combines original features in RoboMap V2 with higher throughput while OpenView is running.

RoboMap now moves OpenView 4.1 submap contexts between maps or between OpenView computers running OpenView 4.1. It can hide objects in the target map that were hidden in the source map. It can also "unmanage" objects in the target map that were unmanaged in the source map. It can hide objects in the target map that did not exist in the source map and delete objects in the target map that were not in the source map. It can rebuild executable icons without using the actionmap mapping file. In addition, RoboMap has a directory browser for backing up and restoring image files, expanding and collapsing directories, and progress meters showing the percentage completed for symbols and submaps. RoboMap also automatically compresses and uncompresses image

Cognos Axiant 4GL Version 2.0



Client-Server Development

Cognos has introduced Axiant 4GL Version 2.0, a visual development environment that leverages investments in PowerHouse applications while taking advantage of the newest development technologies. Also available are PowerHouse Power2000 and a Web-based agent that allow developers to download a small portion of this Year 2000-compliance tool to identify whether there are Year 2000 date issues in PowerHouse applications.

Axiant 4GL features a GUI for increased ease of use; the ability to move to relational databases and newer technologies; support for Windows NT and Windows 95, as well as UNIX, MPE/iX, etc., servers; and deployment to a variety of client-server configurations simultaneously.

PowerHouse Power2000 combines a complete understanding of PowerHouse 4GL language, mass editing capability, and a knowledge base of where Year 2000 issues can affect an application.

Axiant 4GL starts at $4,500 per developer. PowerHouse Power2000 starts at $15,000 per license.

Contact Cognos, phone: (617) 229-6600, http://www.cognos.com.

files to reduce disk space needed. UNIX compress and GNU *gz* files are both supported.

Remote Access for HP OpenView

Onion Peel Software has announced its Remote OpenView Environment (ROVE) for HP OpenView, which allows remote access to OpenView information from PCs, terminals, and palmtops by sharing a single OpenView map.

ROVE enables network troubleshooters to gain access to OpenView information without being at an X-terminal in the office. ROVE deploys a streamlined character interface that runs on slow dial-up lines from inexpensive terminals and is designed to provide all of the significant monitoring capabilities of OpenView to administrators, technicians, and occasional users without the need for expensive LANs and X-terminals.

The product reduces the memory and processor bottlenecks of HP's current X-windows implementation. The two versions—ROVE v100 and ROVE Motif (which can be mixed) allow real-time status, local customized menus, event log access, and full map navigation. Up to 32 clients can be supported simultaneously, with each client requiring less than 5 MB of memory.

ROVE is supported on HP-UX 9.x and 10.x.

Contact Onion Peel Software, phone: (919) 571-7910, fax: (919) 571-8338, e-mail: sales@ops.com, http://www.ops.com/.

Motif Graphing Widget

KL Group, Inc. has announced Release 3.0 of XRT/graph, charting widget for UNIX and OpenVMS user-interface developers.

Version 3.0 includes new graph types,

Make Performance Tuning Easy With SarCheck™

SarCheck *is the answer to your performance tuning problems. Regular use lets you know exactly what's happening with your system resources and helps you maintain control of them.*



Understand what's Happening
SarCheck helps you make sense out of sar data by translating pages of confusing output into a concise, easy-to-read report - complete with tuning and upgrade recommendations.

Maintain Full Control
Regular use of ***SarCheck*** keeps you in control of your system resources and lets you make intelligent, informed decisions about them.

Plan for Future Growth
SarCheck's Built-in capacity planning features allow you to plan for growth - before slow-downs or problems occur.

SCO OK

AURORA SOFTWARE INC.

HEWLETT PACKARD Channel Partner

PO Box 1033, Plaistow, NH 03865
(603) 382-4200 Fax (603) 382-4247 http://www.sarcheck.com

CIRCLE 26 ON READER SERVICE CARD



HP/APOLLO®
Sun Microsystems®
WORKSTATIONS

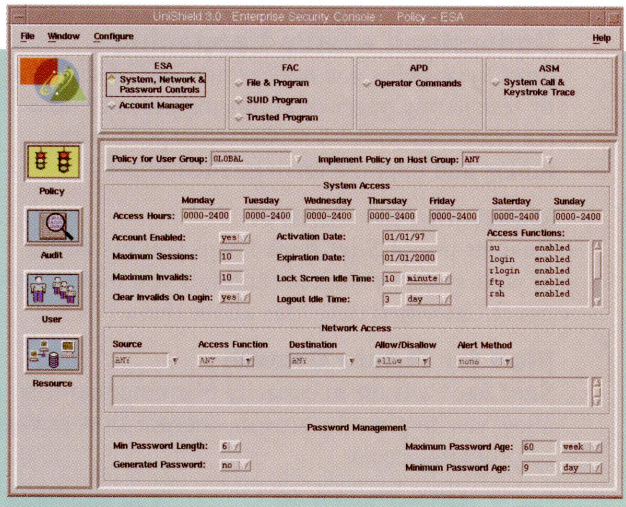
■ Quality depot repairs at substantially reduced cost

On-Site Maintenance for HP/Apollo - Sun Mircrosystems Savings up to 50%	Systems/Spare Part Sales Technical Assistance Buy ■ Sell ■ Repair ■ Service Satisfaction Guaranteed

CALL US TODAY AT
AMC COMPUTER SERVICES, INC.
146-B Rangeway Rd., P.O. Box 286
N. Billerica, MA 01862
508-670-9395 ■ FAX 508-670-9327
1-800-878-2627
E-MAIL: AMC@TIAC.NET

CIRCLE 3 ON READER SERVICE CARD

NIT UniShield 3.0



Intranet Security Software

Network Information Technology (NIT) has announced UniShield 3.0, designed to protect against security incidents, detect security violations, and enforce security policies on corporate intranets.

UniShield 3.0 centralizes security administration for intranets, offering access control to network services and hosts down to the file, program, and operator command level. UniShield does not replace firewalls and encryption but supplies a "final line of defense" against internal and external security threats.

The product provides centralized security administration, password management, granular resource access control, privilege delegation real-time monitoring and alert, and filterable audit reporting. It prevents malicious Java applets from executing unauthorized commands and unknown ActiveX Controls from damaging critical resources. With UniShield 3.0, an administrator can implement file and application access control policies for 500 hosts or more.

Contact NIT, phone: (408) 996-1688, fax: (408) 996-1898, e-mail: info@nit.com, http://www.nit.com/.

property pages, enhanced barcodes, enhanced axis labels, CDE compliance, Motif 2.0 support, and online .PDF documentation.

XRT/graph 3.0 is also available as part of KL Group's XRT Professional Developer's Suite, which includes XRT/graph, XRT/table, XRT/3d, XRT/field, and XRT/gear.

XRT/graph 3.0 is priced at $1,995 in North America and at $2,400 internationally.

Contact KL Group, phone: (800) 663-4723 or (416) 594-1026, fax: (416) 594-1919, e-mail: info@klg.com.

Web Development

Prolifics has announced Prolifics 2, which introduces Web-based report writing, client-side processing, native browser extensions, and intuitive end-user interfaces for the untrained browser user.

Prolifics' ability to perform run-time load balancing assures optimal performance when transaction volume peaks. It also enables customers to initially deploy 2-tier, departmental, and workgroup applications with a smooth migration path to a multitier architecture that can support a larger number of users and more complex requirements.

The product provides rapid report generation, which is automatically generated as HTML and can be invoked interactively or in batch mode. It extends support for automatically generated JavaScript.

Prolifics' two-tier tool, JAM/WEB, is priced at $5,150. Its three-tier tools are priced at $35,000.

Contact Prolifics, phone: (888) 955-7666 or (212) 267-7722, fax: (212) 608-6753, http://www.prolifics.com.

Cross-Platform Development

TakeFive Software, Inc. has announced a new integration module from VERILOG for SNiFF+. The module, VERILOG's ObjectPartner, is a CASE tool based on OMT analysis and design methods. ObjectPartner generates C++ code, which is seamlessly loaded into the SNIFF+ development environment.

The integration of SNiFF+ and ObjectPartner provides streamlined tools for any task in the software life cycle within one single environment. Developers can freely mix and match hardware platforms, compilers, debuggers, editors, third-party tools, and programming languages.

Its Fuzzy-Parser technology makes SNiFF+ powerful and fast. SNiFF+ and all integrated tools feature an intuitive and easy-to-learn GUI. SNiFF+ simplifies development and management of projects with languages like C, C++, Java, FORTRAN, and IDL. Powerful browsing tools make it easy to sort through hundreds of source files. For the testing phase, ObjectPartner automatically takes the project structure issued by SNiFF+ to compute online test coverage measurement and to define test strategies.

Contact TakeFive Software, phone: (408) 777-1440, fax: (408) 777-1444, e-mail: info@takefive.com, http://www.takefive.com/.

Object Request Broker for Java

ObjectSpace, Inc. has released a public beta version of Voyager, said to be the first agent-enhanced object request broker (ORB) for Java. Developers building traditional distributed systems find Voyager much more advanced and easier to use than platforms such as RMI and CORBA, the company notes. Voyager supports object mobility and mobile autonomous agents for building agent-enhanced systems.

Voyager is designed to be used for systems development on general purpose platforms as well as embedded applications. Its small size and flexible design allow it to easily be integrated into consumer electronic devices and used with a wide variety of communications protocols.

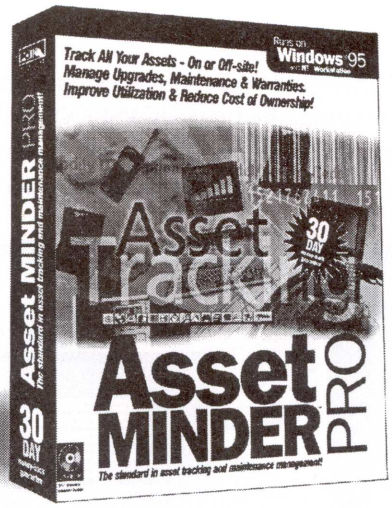
Voyager is now available for free download at http://www.objectspace.com. The production version of Voyager will be available in the third quarter of 1997. In the future, ObjectSpace will offer significant add-on capabilities, a complete suite of support services, and a comprehensive developer support program. As with the beta version, the production Voyager platform will be free for most commercial uses. Source code licenses are also available.

Contact ObjectSpace, phone: (972) 934-2496, fax: (972) 663-9099, http://www.objectspace.com.

Web Conferencing

O'Reilly & Associates has announced *Building Your Own Web Conferences*, which

Losing Track?



If you've lost track of what I.T. assets you have and where they are – you need Asset MINDER PRO. Asset MINDER PRO tells you what you have, where it is and when it needs maintenance.

Bar-Coding lets you track physical locations, changes and off-site loans quickly and easily.

Keep track of versions, models, serial numbers and configurations and use Internet buttons to instantly connect to important web-sites or e-mail addresses.

For a limited time:
$299 standalone version
$149 per additional user
includes one free upgrade!

PC IMAGEWARE CORPORATION

1-888-Imageware 462-4392

email: pci_info@pci-corp.com web: http://www.pci-corp.com

CIRCLE 57 ON READER SERVICE CARD

Hewlett-Packard 9000 It's Our Specialty

200/300 Series	400 Series	700 Series	800 Series
216/236/217	425e	705/710	E/F/G
310/320/330	425t	715/720	H/I/K
350/360/370	425s	730/750	
318/319/340	433s	735/755	
345/375/380		C/J Series	

We also carry memory and interface for all of our workstations

Printers	Mass Storage	Plotters
2225A/B/C/D	9121/9122	DesignJets
3630A PaintJet	9153A/B/C	DraftPro
C1602A PaintJet XL	C2254HA	DraftMaster
LaserJet II/IID/IIP	C2440HA/JA	Desktops
LaserJet III/IIID	C3232A	Electrostatic
LaserJet IIIP,IIIsi		
LaserJet 4L/4P/4,/4+/4SI/4V/4MV	C2213A/D	
	DAT Drives	
Specials	CD ROMs	
DesignJet 650C Plotter	Optical Drives	

We offer large discounts, outstanding service and immediate delivery.
TED DASHER & ASSOCIATES
PH: 800-638-4833
FAX: 205-591-1108
E-mail: sales@dasher.com

HEWLETT PACKARD Channel Partner

CIRCLE 49 ON READER SERVICE CARD

Networked Faxing

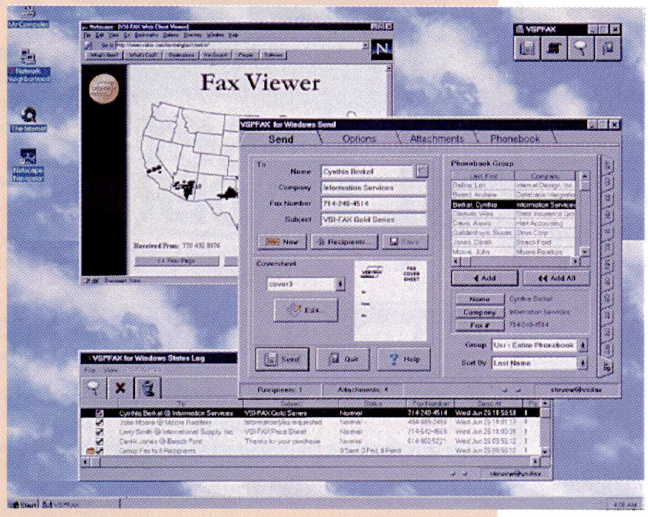
V-Systems, Inc. has announced the VSI-FAX Gold Series, which extends VSI-FAX across UNIX and Windows NT networks and the Internet. VSI-FAX Gold Series supports everyone in the enterprise, with direct-to-the-desktop personal fax capabilities. The product can send invoices, purchase orders, notices, etc. directly to recipients.

The fax server is available on UNIX and Windows NT platforms. The fax client is available on Windows, UNIX, and Web platforms.

Prices of the VSI-FAX Gold Series Fax Servers range from $1,195 to $2,995. The Virtual Fax Server is priced at $795. Network clients are priced from $59.

Contact V-Systems, phone: (714) 489-8778, fax: (714) 489-2486, e-mail: info@vsi.com, http://www.vsifax.com.

V-Systems VSI-FAX Gold Series



offers a Web application for creating real-time, Web-based discussion groups. The book/CD package includes O'Reilly's WebBoard 2.0 Web conferencing software for Windows 95 and NT.

With WebBoard 2.0, any Web site can host live, real-time threaded chat discussions on multiple topics with unlimited participants. The conference administrator can display ads or messages in chat rooms. The software can be used with WebBoard's own multithreaded internal Web server or with a Windows CGI- or ISAPI-compliant Web server.

The version of WebBoard 2.0 in the *Building Your Own Web Conferences* book/CD supports two virtual boards with 10 conferences per board and retails for $59.95. WebBoard 2.0 XL, an extended-license version of WebBoard, lets Web administrators create up to 255 virtual boards, with an unlimited number of conferences in each virtual board. *Building Your Own Web Conferences* (ISBN: 1-56592-279-4) and the WebBoard 2.0 XL license are available as a package (the WebBoard Power Pack) directly from O'Reilly & Associates for $399.

Contact O'Reilly & Associates, phone: (800) 998-9938 or (707) 829-0515, fax: (707) 829-0104, e-mail: info@ora.com.

Math Solutions

Solveware, Inc. has developed MATHANSR, a program to make math solutions easy for the expert as well as the novice. In solving right triangles, for instance, only two numbers are needed to obtain all the dimensions for a right triangle. There are separate solutions for each of many math problems encountered in the real world. The solutions are displayed with graphics as well as tables.

MATHANSR was developed to fill a need for obtaining accurate answers, quickly, on the shop floor of a mold making business. It is a spreadsheet program that runs on Microsoft's Excel. The user can only enter numbers in the "unlocked cells," making the program very forgiving. By using some of the features of Excel, any answer can be obtained for any dimension on the graphically displayed figures. Colors are used extensively to aid in communication.

Contact Solveware, Inc., phone (561) 575-4502, fax: (561) 575-9841, http://www.solveware.com/.

Java-Based Data Analysis

Visual Numerics, Inc. has announced JWAVE 1.0, which enables PV-WAVE users to create visual and numerical data analysis client applications written in Java. These Java applications communicate with a PV-WAVE server, which acts as a graphics and numerics engine for the client.

JWAVE provides Java classes for plotting data and supports the full range of PV-WAVE capabilities. At the same time, it bars access to critical data.

JWAVE consists of a set of Java components (classes) based on Netscape Internet Foundation Classes (IFC). Application developers can create Java programs that perform complex numerical analyses and generate 2-D plots, surfaces, contours, animations, and other graphical data representations. JWAVE supports real-time collaboration and provides necessary encryption features.

A free evaluation copy of JWAVE can be downloaded from Visual Numerics' Web sit at http://www.vni.com. Its start-

ing cost is $3,395.

Contact Visual Numerics, Inc., phone: (800) 364-8880 or (303) 530-9000, fax: (303) 530-9329, e-mail: margaret@boulder.vni.com.

Project Management

Micro Planning International (MPI) has announced MICRO PLANNER X-Pert - Version 2.4.2, and MICRO PLANNER Manager - Version 1.4.2. The latest releases offer added features and enhancements, including an automatic update feature, "Year 2000" compliance, and increased speed.

MICRO PLANNER can be used independently or can be combined to provide a total enterprise-wide solution. MICRO PLANNER Manager is ideal at the project level, while X-Pert is designed for use at higher levels in the organization for consolidating and managing many projects or subprojects.

MICRO PLANNER also accommodates multiple platforms by offering complete "file open" compatibility between Macintosh, UNIX, and Windows versions. MICRO PLANNER X-Pert is priced at $1,995. MICRO PLANNER Manager is priced at $695.

Contact MPI, phone: (800) 852-PLAN or (303) 757-2216, fax: (303) 757-2047.

Mainframe-Class Storage

Zitel Corporation has announced 9-GB drives in its CASD-II/Enterprise storage solution. CASD-II/Enterprise features "intelligence-based" caching technology, mainframe-class design, and modular architecture to deliver performance, scalability, reliability, availability, and serviceability. The CASD-II/Enterprise entry-level storage module system includes 36 GB of disk, 64 MB (scalable to 512 MB) of ECC

Make the UNIX to MPE Connection

IX/92™

Full featured HP terminal emulation

New! Version 6

Faster File Transfer
NS/VT Network Option
Enhanced Script Language

Available for:

HP-UX Interactive UNIX SCO UNIX SunOS/Solaris

Software Licensing Corp., Suite 280, 930 Tahoe Blvd. Unit #802
Incline Village, NV 89451-9436
Phone: (800) 831-0882 or (702) 832-0881 Fax: (702) 832-0883
All trademarks are the property of their respective holders.

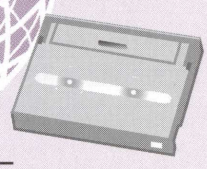
CIRCLE 180 ON READER SERVICE CARD

Large Databases Require New Solutions for Backup !!!!!

FINALLY! Backup 4 to 8 Gigabytes/Hour/Drive
Store 20 to 40 Gigabytes/Tape
Random Libraries from 5 to 264 Cartridges

DallasTools

Back-up, Restore and Archiving Software for Heterogeneous Networks



Dallastone also sells and supports all major brands – Quantum, Breece Hill, ADIC and Odetics

DALLASTONE
2 Cote Lane
Bedford, NH 03110

Phone 603-647-8168
Fax 603-624-2466
Email dtool@delphi.com

CIRCLE 76 ON READER SERVICE CARD

protected cache, Dual Port/Dual Path SCSI interface, redundant power system, and a combination battery charger/power supply and rechargeable battery. The N+1 "Plus" power system protects the data in cache memory (non-volatile RAM or NVRAM) by safely destaging the data to disk in the event of a power failure. Additionally, each time another storage module is added to a configuration, capacity, performance, and connectivity are scaled linearly.

Contact Zitel, phone: (510) 440-9600 or (800) 622-5020, fax: (510) 440-9666, http://www.zitel.com.

Order Fulfillment Solution

Industri-Matermatik International Corporation (IMI) has announced Version 5 of its System ESS demand chain management solution. System ESS helps companies maintain timely and accurate information flows across global order fulfillment planning and execution processes.

Version 5 features an object-oriented, standards-based architecture that enables users to tailor System ESS. It also provides a "plug-and-play" environment, enhanced integration with other solutions (Oracle, Manugistics, and Metasys), and support for Java-based development.

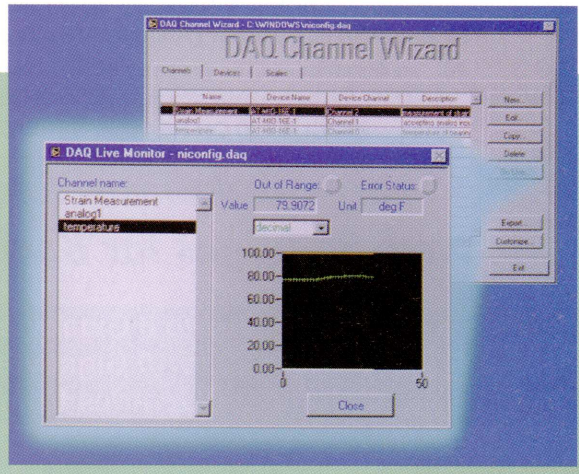
A new service management component addresses the needs of vertical markets, such as wholesale/retail. Redesigned electronic commerce components make it easier to conduct business using EDI and Internet messaging. The Business Engineering Workbench documents changes to System ESS and provides a graphical online help facility.

Pricing is based on customer configuration and typically starts $600,000. System ESS runs on HP-UX and other platforms.

Contact IMI, phone: (914) 631-2700, fax: (914) 631-5111, http://www.im.se.

Graphical Instrumentation

National Instruments has announced Version 4.1 of its LabVIEW graphical instrumentation software. LabVIEW is a data acquisition (DAQ) and analysis software package that streamlines DAQ development through automatic, point-and-click program generation. Users simply describe their measurement connections, and LabVIEW generates a ready-to-run program that meets their specifications. It can complete this entire process in only a few minutes, the company notes. The generated programs are native LabVIEW "G" source code.



National Instruments LabVIEW Version 4.1

Other enhancements include an online tutorial for Windows users, automated links to important LabVIEW sites on the Web, and an Object Linking and an OLE automation interface for automated report generation via HiQ for Windows 95 and NT. LabVIEW 4.1 is available on Windows NT/95/3.1, HP workstations, and other platforms.

LabVIEW pricing starts at $995. Existing LabVIEW for Windows users can upgrade to Version 4.1 for $295.

Contact National Instruments, phone: (800) 258-7022 or (512) 794-0100, fax: (512) 794-8411, e-mail: info@natinst.com.

AWK Programming Guide

Specialized Systems Consultants, Inc. (SSC) has announced *Effective AWK Programming, Second Edition* by Arnold Robbins, a user's guide for AWK and its GNU implementation, *gawk*. Information is presented in a tutorial format with practical tips and examples presented in a clear and concise manner. The appendices include a complete language summary, installation information, and a glossary. This edition adds updates for GNU *gawk* version 3.0.3 and is bound with SSC's *AWK Reference* card, also by Robbins.

AWK, a reporting language based on the concepts of the UNIX tools *grep* and *sed*, is a sophisticated language that includes user-defined functions and dynamic regular expressions. AWK is defined in the POSIX Command Language and utilities standard and is available for many platforms, including MS-DOS, VMS, and most UNIX platforms.

Effective AWK Programming is priced at $27.

Contact Specialized Systems Consultants, Inc., phone: (206) 782-7733, fax: (206) 782-7191, e-mail: sales@ssc.com, http://www.ssc.com/.

Ultra SCSI Interconnection

ATTO Technology announced SCSI Sidekick, an Ultra SCSI operating system independent device that combines the benefits of a SCSI signal converter, bus extender, and bus isolator into one unit.

The SCSI Sidekick converts the SCSI signal, enabling a single host adapter to communicate with a combination of single-ended and differential SCSI devices on the same SCSI bus with no performance degradation. Integrators can now interface the noise reduction and extended cabling advantages of popular differential SCSI devices with the vast base of low cost, single-ended products without the need for system reconfiguration. A bus isolation switch allows workgroup members to disconnect the device from the host without shutting down the entire system.

SCSI Sidekick's operation is transparent and does not occupy a SCSI ID. In addition, the unit provides Upper-Byte Termination, which allows communication with both narrow (8-bit) and wide (16-bit) devices on the same bus.

ATTO SCSI Sidekick is priced at $295.

Contact ATTO Technology, Inc., phone: (716) 691-1999, fax: (716) 691-9353, e-mail: mkt@attotech.com, http://www.attotech.com.

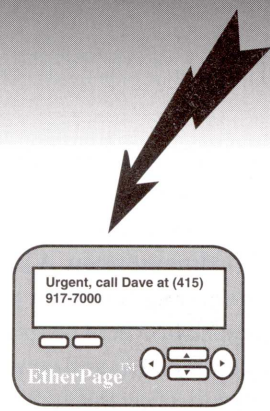
Print Spooler

LBM Systems has announced a Java Network Client for UniQue Print, LBM's mainframe-caliber print spooler for UNIX.

In many network environments, a single UNIX machine functions as a print

ALPHANUMERIC PAGING FOR UNIX

ROBUST, RELIABLE, USER-FRIENDLY DELIVERY OF MESSAGES ANYTIME, ANYWHERE



- Email forwarded to pager automatically
- Pages can be generated from scripts and network monitoring programs
- GUI and command line interface
- Works with any paging service
- Automatic email confirmation, history logs and error reporting
- Client-server technology
- Works with digital and alphanumeric pagers

Personal Productivity Tools
for the Unix Desktop

14141 Miranda Rd
Los Altos Hills, CA 94022
Email: sales@ppt.com
Tel: (415) 917-7000
Fax: (415) 917-7010
http://www.ppt.com

CIRCLE 91 ON READER SERVICE CARD

Don't Gamble with Your Data.. Cover Your Assets!

The Model 4/30 DLT Tape Library, offered by Concorde Technologies, provides state-of-the-art automated backup and archiving to small- and medium-sized departments and workgroups. The 4/30 is easily scalable, featuring the lightest and most compact robotics ever designed, transports tapes faster and more efficiently than linear robotics systems, has transfer rates of up to 144 gigabytes per hour and has an average access time of less that 6 seconds.

with state-of-the-art automated backup and archiving

Concorde Technologies offers a selection of tape systems that provide the fastest access time in the industry. These tape-base systems, of which the 4/30 is the newest and smallest, give users access to large data files and quick retrieval times through industry-leading robotic designs.

Concorde Technologies is a leading provider of optimized storage solutions for Open System users. The company's products and services are used to store, transport and secure more than 100 petabytes of information, ranging from mainframe data to client/server applications to video and still images.

CONCORDE TECHNOLOGIES
BOOTH 2009

Concorde Technologies, an HP Channel Partner and a StorageTek Global Partner

server, which allows ease of administration and trouble-shooting, centralized control of system resources, and simplified printer administration. A Java manager program, PC/NT client printing software, and UNIX client programs provide complete access to and control over all printing functions through UniQue Print without having to log on to the UNIX print server. A single GUI reference screen with color-coded status displays allows printing output to be monitored at a glance. Without logging onto the server, remote users have the ability to start and stop jobs; print partial jobs; change print jobs according to form, class, group, number of copies, page range, and queue position; schedule printers by class and/or group; control paper and stock by on screen prompt; access any filter, forms overlay, fax gateway, or virtual printer; and handle remote and local printers identically.

Contact LBM Systems, phone: (203) 966-0661, fax: (203) 966-8242, http://www.lbmsys.com.

SCSI Connectivity

APCON, Inc. has announced its newest SCSI Bus Extender, PowerLink. PowerLink extends the SCSI bus to 10 km while supporting Ultra Wide SCSI performance at a maximum SCSI data rate of 40 megabytes per second.

PowerLink provides real-time access to remote peripheral devices such as RAID arrays, tape libraries, optical juke-boxes, and CD-ROM towers that are located up to 10 km from the host computer. PowerLink supports low-cost coax for extensions up to 30 meters, high-performance coax to 200 meters, a multimode fiber optic interface for distances to 550 meters, or optionally, a single-mode optics interface to 10 km.

A GUI allows the system administrator to remotely operate and control selection of extension interfaces, report SCSI bus activity, gather performance statistics, and initiate on-board diagnostics.

The Powerlink external unit, Model ACI-2030 retails for $4,000 for the single ended version and $4,500 for differential. Both versions are available off the shelf. Lower-cost board versions are also available for complete internal system integration.

Contact APCON, phone: (503) 639-6700, fax: (503) 639-6740, e-mail: info@apcon.com.

Decision Support Tool

Leading Market Technologies, Inc. (LMT) has announced EXPO/Econometrics, an add-on module to LMT's EXPO analytics and decision support platform.

EXPO/Econometrics provides "point-and-click" access to a wide range of tests and estimation techniques; integration with all major data distribution platforms and major market data sources, including OPEN BLOOMBERG, Reuters, Bridge, FAME, and LIM; and important innovations in applied econometrics requested by LMT's clients.

The EXPO/Econometrics module provides a full set of econometric tests, transformation functions, correlations, and estimation techniques including ARCH, rolling regression, multivariate regression, seasonal adjustment with moving averages, post-regression residuals tests, and multicollinearity testing.

The EXPO/Econometrics module is priced at $3,995 for Windows 95/NT and all major UNIX platforms, including HP-UX.

Contact Leading Market Technologies, phone: (617) 494-4747 ext. 235, fax: (617) 494-4788, e-mail: sgalla@world.std.com.

Object-Relational Application Development

Black & White Software announced D/Enable Version 2.0. DB/Enable is comprised of CORBA-based adapters that encapsulate relational databases as object-based servers, supporting the access and retrieval of arbitrary data structures. DB/Enable provides extensive graphical development tools to interactively create, modify, test, and generate code for Oracle, Sybase, and Informix applications.

CORBA employs IIOP as the basis for network communication among heterogeneous components and uses a drag-and-drop metaphor for easily attaching database queries and responses to user interface components. Source code is automatically generated that utilized the DB/Enable convenience library.

New features include a C++ interpreter for testing code on the fly, three-tier architecture, the ability to encapsulate relational database queries into reusable objects, and support for simultaneous connections to multiple dissimilar databases.

DB/Enable 2.0 is available now as a beta release on Solaris-2 and HP-UX 10.x. DB/Enable is priced at $3,000.

Contact Black & White Software, phone: (408) 369-7400, fax: (408) 369-7406, e-mail: info@blackwhite.com.

C++ Graphics Toolkit

Interactive Network Technologies, Inc. (INT) has announced Version 1.1 Carnac, its 2-D graphics toolkit. Two

Software Developers Kits (SDKs) have been added to the toolkit, the Flavors SDK and Dynamic Scenes SDK. Carnac's extensibility allows developers to tailor the product to their exact requirements.

The Flavors SDK simplifies the production of user-defined primitives (complex shapes), while the Dynamic Scenes SDK assists with the production of the code for generating dynamic scenes on the fly. In addition, support for 1, 8, and 24-bit rotatable, shareable, and scalable images with full transparency control has been added. Finally, CGM output has been added to the existing options, which already include support for a wide range of printers and PostScript format.

Carnac is available for Windows 95, Windows NT, and most common UNIX/X Window implementations.

Contact INT, phone: (713) 975-7434, fax: (713) 975-1120, e-mail: info@int.com, http://www.int.com. ◧

Attention vendors: New product announcements should be sent to New Products Editor, hp-ux/usr magazine, *Interex, P.O. Box 3439, Sunnyvale, California 94088-3439, USA, or e-mail: pollace@interex.org.*

Deadline for submission is two months prior to publication.

Do You Know Where Your Security Holes Are?
Find Them with SecurityAudit/UX!

Have You Heard that UNIX is Notorious for Its Lack of Security Features?
Do You Know Where to Check to See if Your HP-UX System is Secure?
Do You Have the Time to Do This Checking Regularly?
Use SecurityAudit/UX To Do It All!

EVEN IF YOUR SYSTEM IS SET UP CORRECTLY (AND HOW WOULD YOU KNOW IF IT WAS?) AND HAS NO SECURITY LOOPHOLES, IT CAN BE VERY DIFFICULT TO MONITOR SYSTEM CHANGES, AND TO ENSURE THAT SECURITY ISN'T COMPROMISED. THE MAGNITUDE OF THE PROBLEM INCREASES AS THE TOTAL NUMBER OF USERS CONFIGURED AND THE TOTAL NUMBER OF FILES GROWS. IT'S EASY FOR ORDINARY USERS TO CHANGE THE SECURITY OF THEIR OWN FILES TO ALLOW OTHERS TO ACCESS THE CONTENTS. A LOOPHOLE LEFT BEHIND INADVERTENTLY OR ON PURPOSE MAY BE EXPLOITED BY A DISGRUNTLED EMPLOYEE OR A HACKER TO BREAK SYSTEM SECURITY, SOMETIMES MUCH LATER.

SecurityAudit/UX PRODUCES OVER 40 REPORTS, CONTAINING DETAILED INFORMATION ON THE FOLLOWING CLASSES OF PROBLEMS:

- **User and Group-related problems, including weak passwords and non-unique identification numbers.**
- **File-system related problems, including historical tracking of files and detection of potential Trojan horses.**
- **PDF-related security problems, extended to detect changes in ACL specifications.**
- **Logging subsystems status display, and logfile analysis.**
- **Network-related status display and configuration weaknesses.**

SecurityAudit/UX RUNS ON ALL HP-UX BASED 9000 SERIES 700 AND 800 SYSTEMS, AND HAS BEEN SPECIFICALLY TAILORED TO ADDRESS PECULIARITIES OF HP-UX, SUCH AS PDF, ACL AND HP'S SHADOW PASSWORDS.

Call EUGENE VOLOKH for more info!

VESOFT
Purveyors of Fine HP Software
Since 1980
(310) 282-0420

1135 S. Beverly Drive
Los Angeles, CA 90035 U.S.A.
FAX (310) 785-9566

CIRCLE 40 ON READER SERVICE CARD

Get On Board!



Computer Express est. 1983
... The Fast Track to Success

HP3000 Professionals

Immediate openings in New England, Florida, Kentucky and Texas

Established in 1983, Computer Express is a nationally known I.T. consulting firm based in Boston (Wakefield, MA). We have immediate openings for skilled HP3000 professionals with experience in:

- **COBOL**
- **POWERHOUSE**
- **ASK**
- **SPEEDWARE /PROTOS**

If you are experienced in any of the above areas, we can offer competitive wages, Section 125 Medical/Dental, 401(k), Tuition Reimbursement, Vacation/Sick time, and more! Contact our Recruiting Department for more details.

CEI

Computer Express, Inc.
301 North Ave., P.O. Box 308, Wakefield, MA 01880
Voice: 617-246-4477, Fax: 617-246-1434
E-Mail: jobs@ComputerExpress.Com
Web: http://www.ComputerExpress.Com

EOE - Member NACCB

CIRCLE 34 ON READER SERVICE CARD

Sign up NOW for Spring 1998 Listings

hp-ux/resource directory

The *hp-ux/resource directory* is a complete resource guide for HP-UX users seeking answers. This is one of the industry's most extensive reference guides for HP-UX products, services, and vendors. It will be devoted entirely to HP 9000 users operating in multi-user, workstation, and multi-system UNIX environments. This bi-annual directory, published each year in March and September, is a separate publication mailed out with *hp-ux/usr* magazine, the only HP-specific publication on the market.

Added BONUS: your message will reach your customers for one full year on the *Internet.* Look for the directory on the Interex home page http://www.interex.org. The investment for a full year listing in the *hp-ux/resource directory* is $475.

Product Categories

Accounting
Accounting/Financial
Accounting Software
Alphanumeric Paging Software
Application Development Software
Application Development Tools
Application Development Tools/4GL
Application Engineering
Automated Software Quality
Backup Hardware
Backup/Restore
Backup Software
Bar Code Data Collection Systems
Batch Job Management
Books
Business-Critical Application Development & Deployment
Business Software
CD-R
CAD Software/Hardware
Change Management for Software Development
Change Management Tools
CheckPoint Restart Facility
Client-Server
Client-Server Software
Communications
Communications Servers
Communications Software
Consulting
Consulting/Systems Integration
Customer Support
Customer Support/Help Desk Systems
Database Management Systems
Database Management Tools
Data Center Management
Data Migration Tool
Data Warehousing
Decision Support Systems
Disaster Recovery
Distributed Computing
Distribution Software
Distributor
Document Management
Editors
Electronic Data Interchange (edi)
Electronic Form Printing
E-Mail & Directory Integration
End-User Access Tools
End-User Computing
Equipment
Facility Maintenance Software
Executive Information Systems
Fault Management
Fax Automation
Fax Software
File Manager Utility
Financial
Fourth Generation Language
GIS (Geographic Information System)
Government & Utility Software
Graphics
Groupware
Hardware
Hardware/Mass Storage
Help Desk Management
Hardware Subsystems
Hierarchical Storage Management
Human Resources & Personnel Systems
Image Processing
Image Storage & Retrieval Management
Industrial Terminals
Instrument Control
Integration Tools
Interface
Internet
Internet Commerce
Internet/Intranet
Internet Services
Internet Solutions
Inventory Control
I/O Boards
Job Scheduling & Workload Management
Justice Software
Laser Printing Software
Maintenance
Manufacturing Software
Mass Storage
Mass Storage Peripherals
Math Library
Memory
Memory Upgrades
Middleware
Migration Services
Migration Services/Tools
Migration Tools
Millennium Solutions
Multimedia
Network Backup Software
Networking
Networking Systems
Network Management
Output Management
Payroll
PC Card Reader
PC Integration
Performance
Performance Software
Personnel Management
Pointing Devices
Personnel Services
Power Protection & Conditioning
Printers
Print Management
Print Management Software
Process Control Software
Production Planning
Project Management
Protocol Converters/Interfaces–Hardware
Publications
Public Safety Software
Quality Assurance Tools
RAID
RAID Disk Arrays
RAID Storage
Records Management
Remote Networking
Rentals
Report Viewing, Printing, & Distribution
Report Writers
Sales & Marketing
Scheduling
Scheduling/Task Management
Security
Service Repairs
Software
Software Backup
Software Development Tools
Software Distribution Tools
Software Maintenance & Testing
Spoolers
Spreadsheets
Statistics/Data Analysis
System Integration
System Management
System Management Tools
System Printers
3-D Graphics Tool Kit
3-D Porting Tool
Tape Backup Products
Tape Storage/Data Interchange
Technical Documentation/Cross Referencing
Terminal Emulation
Text & Information Retrieval
Time & Billing
Time Reporting Terminals
Training
User Groups
Video/Keyboard/Mouse Extension
Warehouse & Distribution Management
Workstations

Other categories may be created as needed.



interex

Shared Success for HP Computing Professionals

hp-ux/resource directory

Spring 1998 Listing Form

	FIRST LISTING	EACH ADDITIONAL LISTING	TOTAL
☐ Listing (includes two issues)	$475	$375	$_____
☐ Hyperlink to your home page	$500	$150	$_____
☐ $1.00 Per Word Over 75 Words			$_____
☐ Company Product Logo or Photo	$100	$ 50	$_____
☐ Cross Reference	$200	$200	$_____
		Total	**$**_____

Closing Date: Friday, December 12, 1997

1. _____ Category _____ Product Name
 _____ Operating Environment
2. _____ Category _____ Product Name
 _____ Operating Environment
3. _____ Category _____ Product Name
 _____ Operating Environment
4. _____ Category _____ Product Name
 _____ Operating Environment

Product Description

Please attach product description. Be sure to indicate product name and operating environment above.

Note: There is a $1.00 per word charge for each listing over 75 words.

Company Name _____

Address _____

City _____ State _____ Zip _____

Telephone _____ Fax _____

Web URL _____ E-mail: _____

Authorization:

Signature _____ Print Name _____

Title _____ Date _____

INTER-OFFICE USE ONLY:

NEW LISTING: (*never appeared in either directory in any category*)
LOGO: ☐ PHOTO: ☐ LINK: ☐

RE-NEW LISTING: (*new to category and/or directory*)
COPY: ☐ New ☐ Changes ☐ Pickup
LOGO: ☐ New ☐ Pickup ☐ None
PHOTO: ☐ New ☐ Pickup ☐ None
LINK: ☐

REPLACE LISTING: (*existing copy in hprd, in same category*)
COPY: ☐ New ☐ Changes ☐ Pickup
LOGO: ☐ New ☐ Pickup ☐ None
PHOTO: ☐ New ☐ Pickup ☐ None
LINK: ☐

Please mail or fax completed form and listings to:

Interex, 1192 Borregas Avenue, Sunnyvale, CA 94088, U.S.A., Attention: Kathie Schwartz, 800.468.3739, ext. 620, 408.747.4620, **Fax 408.747.0947**, E-mail: schwartz@interex.org.

PAYMENT OR PURCHASE ORDER MUST ACCOMPANY ALL ORDERS

NOTES:

interex

Shared Success for HP Computing Professionals

BECOME A HOT SHOT HP COMPUTER USER.

GET EXCITING NEW MEMBER BENEFITS!

Interex, the International Association of Hewlett-Packard Computing Professionals, can turn you into a hot shot HP computer user. By joining Interex now, you'll not only receive all the traditional value of Interex membership, you'll also receive exciting, new member benefits!

Your new member benefits will include significant discounts on both technical publications, technical training, and special advisory services.

The technical publication discount gives you a 20% discount on all Prentice Hall and O'Reilly publications. Similarly, the new technical education discount of up to 15% gives members a chance to get the training they need to tackle today's critical computing issues. And through Gartner Group services, you can gain special industry-focused analysis at 20% off.

Get valuable computing solutions throughout the year!

It's easy! Just choose the membership level (see below) that best suits your needs and either call an Interex Membership Rep at 1-800-INTEREX or sign up via the Web at http://www.interex.org/mem.html

Member Benefits:

	Associate $49.50	Contributing $115.00	Contributing plus Online (software) Service $595.00
$200 OFF registration for HP World Conference & Expo	■	■	■
Membership in your **local user group** at discounted rates (if applicable)	■	■	■
Hot News weekly e-mail newsletter	■	■	■
1 Free subscription to ***hp-ux/usr magazine*** (HP 9000 focus)*	■	■	■
1 Free subscription to ***Interact magazine*** (HP 3000 focus)*		■	■
Download Interex publications with keyword search capabilities	■	■	■
Free bi-annual issues of ***hp-ux/resource directory*** (HP 9000 focus)*	■	■	■
Free bi-annual issues of ***Vendor Service Source Directory*** (HP 3000 focus)*		■	■
NEW **20% Discount** on all Prentice-Hall and O'Reilly technical publications	■	■	■
1 Free subscription to ***InterexPress*** (HP, industry and Interex news magazine)		■	■
NEW **Up to 15% Technical education discounts** on NT, UNIX, Oracle, Internet, Java, etc. through ClassPass, Datatech Institute, and Abtech		■	■
Vote in **Interex Board elections**		■	■
Annual Advocacy Survey participation		■	■
System Improvement Ballot participation		■	■
Access **Special Interest Groups** (SIGs)		■	■
NEW **MPE Freeware Tape**—MPE public domain software and utilities			■
NEW **HP-UX 10.x FastStart Toolbox CD**—immediate 10.x functionality			■
Web site **search** and **downloads** from the Contributed Software Library (HP-UX, MPE, RTE) of over 4,800 programs			■
NEW **20% Discount** on Gartner Group Premium research and analysis			■
Contributed Software Library release (HP-UX, MPE, RTE) in the format of your choice			■

**Associate Level Membership is entitled to* hp-ux/usr *magazine OR* Interact *magazine and the* hp-ux/resource directory *OR the* Vendor Service Source Directory.

Membership Application

interex

Shared Success for HP Computing Professionals

The International Association of Hewlett-Packard Computing Professionals

SIGNATURE MONTH / DAY / YEAR

NAME (LAST/FIRST/MIDDLE INITIAL)

TITLE

COMPANY

ADDRESS MAIL STOP

CITY/TOWN STATE/PROVINCE ZIP COUNTRY

TELEPHONE/EXTENSION FAX

E-MAIL ADDRESS

Membership Level (please check one)

- ❑ Associate Level with annual subscription to: ❑ *hp-ux/usr or* ❑ *Interact* **$49.50**
- ❑ Contributing Level **$115.00**
- ❑ Contributing Level Plus Online Service **$595.00**

Preferred CSL Format: (Choose one) ❑ DAT 4mm ❑ 1600 BPI Magnetic Tape ❑ 6250 BPI Magnetic Tape

Additional Software Selections (Choose one or both) ❑ HP-UX 10.x FastStart Toolbox *and/or* ❑ MPE Freeware Tape

Member Activities (open to all levels of membership)

I would like to join and/or participate in the following, please send more information about:

- ❑ Local User Group
- ❑ InterWorks Technical Users Forum of Interex (for HP-UX related system administration, network management and software development)
- ❑ Special Interest Group(s) regarding: _____
- ❑ Presenting at a conference or local user group meeting
- ❑ Writing an article for *hp-ux/usr, Interact* or *InterexPress*
- ❑ Interex Conference and educational events:
 - ❑ HP World
 - ❑ InterWorks
 - ❑ MPE Programmers' Forum (IPROF)
 - ❑ Regional Seminars

Payment Information

- ❑ Check enclosed for $_____,* payable to Interex
- ❑ Bill me. Purchase Order enclosed. PO#_____
- ❑ Charge my: ❑ VISA ❑ MasterCard ❑ American Express

Account# _____ Exp. Date_____

Signature_____

*Canada and Mexico, add $25. All other countries outside the U.S., add $50 for additional postage. Foreign currency is accepted, but payment must be the equivalent of U.S. currency.

Demographics (required)

Your job function:

(check up to three if applicable)

- ☐ Owner/President/CEO
- ☐ VP/Corporate Officer/CFO/Controller
- ☐ CIO/MIS Director
- ☐ Sales/Account Representative
- ☐ MIS Manager
- ☐ Operations/Data Center Manager
- ☐ Technical Manager
- ☐ Scientist/Researcher
- ☐ Electrical Engineer/Mechanical Engineer
- ☐ System Administrator
- ☐ Database Administrator
- ☐ Network Administrator/Manager
- ☐ Systems Analyst
- ☐ Software Developer/Programmer
- ☐ Operator
- ☐ Consultant/Educator
- ☐ Other_____ Please Specify

Primary type of business:

(check one)

- ☐ Manufacturing (computer hardware)
- ☐ Manufacturing (non-computer products)
- ☐ Systems Integrator/VAR
- ☐ Software Developer
- ☐ Finance/Banking/Insurance/Real Estate
- ☐ Retail/Wholesale/Distribution
- ☐ Utilities/Communication Services/ Transportation
- ☐ Government/Military
- ☐ Health Services
- ☐ Education
- ☐ Agriculture/Mining/Construction
- ☐ R & D /Testing/Evaluation Labs
- ☐ Other_____ Please Specify

Service Agreement

If you relocate, should services transfer with you?

☐ Yes ☐ No

Purchases you influence, specify or recommend:

(check all that apply)

- ☐ Computer/Workstation Systems (including Desktops, micros, minis, mainframes, etc.)
- ☐ Peripherals (including printers, disk/tape drivers, modems, memory boards, graphic enhancements, etc.)
- ☐ Software (including application packages, RDBMS, graphic toolkits, communication software, compilers, languages, etc.)
- ☐ Networks (including LANs, WANs, terminal servers, etc.)
- ☐ Services (including on-line services, maintenance, education, training, etc.)
- ☐ None of the above

Application Support

The primary application that you support is:

- ☐ Commercial
- ☐ Technical
- ☐ Both Commercial and Technical

Annual budget

What is your total external IS spending on hardware, software, networking, and contract services for the full current budget year?

	Within your scope	For entire organization
Less than $100,000	☐	☐
$100,000 to $250,000	☐	☐
$250,000 to $500,000	☐	☐
$500,000 to $1,000,000	☐	☐
$1 million to $5 million	☐	☐
$5 million to $10 million	☐	☐
$10 million to $25 million	☐	☐
$25 million to $50 million	☐	☐
$50 million to $100 million	☐	☐
$100 million to $250 million	☐	☐
$250 million to $500 million	☐	☐
Over $500 million	☐	☐

Systems

Computer technologies you are involved with and/or you plan to purchase:

(check all that apply)

	Involved with	Plan to purchase
Unix-based desktops		
HP	☐	☐
Sun	☐	☐
IBM	☐	☐
SGI	☐	☐
Other	☐	☐
Please specify:		
Unix-based servers		
HP	☐	☐
Sun	☐	☐
IBM	☐	☐
SGI	☐	☐
Other	☐	☐
Please specify:		
HP 3000	☐	☐
IBM AS/400	☐	☐
Mainframes	☐	☐
HP Convex	☐	☐
Other Super Computers	☐	☐
HP NetServer	☐	☐
Compaq network servers	☐	☐
Other Intel-based network servers	☐	☐
HP PCs	☐	☐
Other PCs	☐	☐
Other HP Computers	☐	☐
Network Computers	☐	☐
Windows NT desktops	☐	☐
Windows NT servers	☐	☐

Mailing List

Would you like to receive mailings about other computer-related vendors' products and services?

☐ Yes ☐ No

How to Contact Interex:

WWW Complete and submit an electronic application at http://www.interex.org/cgi/memapp/pl

Fax Fax completed form (credit cards only) to: 1.408.747.0947

Mail Mail completed form with payment (if applicable) to: Interex, P.O. Box 3439, Sunnyvale, CA 94088-3439

QUESTIONS?

Call 1.800.INTEREX or 1.408.747.0227

USR897

Advertiser's Index

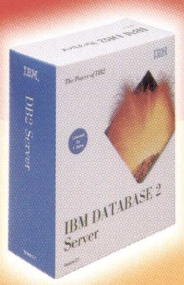
Please call or fill out adjacent card for further product information.

READER SERVICE NUMBER	ADVERTISER	PAGE #
3	**AMC Computer Services** (508) 670-9395/Fax (508) 670-9327	91
35	**APPX** (804) 330-1140	79
2	**Attachmate** (800) 426-6283 or (206) 644-4010	C-4
26	**Aurora Software** (603) 382-4200/FAX (603) 382-4247	91
4	**Bering** (800) 237-4641 or (408) 364-6500	25
145,146, 147	**BMC Software** (800) 811-6766	11,13,15
34	**Computer Express** (617) 246-4477/Fax (617) 246-1434	99
37	**Confluent** (800) 780-2838 or email: info@confluent.com	65
31	**Conley Corporation** (212) 682-0162/Fax: (212) 682-0071	9
76	**Dallastone** (603) 647-8168/Fax: (603) 624-2466	95
44	**Design 3000 Plus, Inc.** (503) 585-0512/ Fax: (503) 585-1706	87
158,159, 160	**IBM** www.software.ibm.com/hp	33,C-2,C-3
33	**IEM** (970) 221-3005 or (800) 321-4671	7
54	**Inotech** (800) 466-8324 or: (703) 641-0469	2
51	**Inclination** (702) 831-5595/Fax: (702) 831-4979	81
55	**Intellisoft** (508) 635-9070/Fax: (508) 635-9210	4
45	**Lotus** http://www.hp.com or http://www.lotus.com	85
58,62	**Lund Performance Solutions** (541) 926-3800/Fax: (541) 926-7723	17,19
96	**MiniSoft** (800) 682-0200	73

READER SERVICE NUMBER	ADVERTISER	PAGE #
38	**Mitchell Humphrey & Co.** (800) 237-0028/Fax: (314) 991-5288	23
7	**Monterey Bay Communications** (408) 429-6144/Fax: (408) 429-1918	1
72	**ORBIT Software** (800) 6-Online or (510) 837-4143	37
57	**PC Imageware Corporation** (888) IMAGEWARE	93
91	**Personal Productivity Tools** (415) 917-7000/Fax: (415) 917-7010	97
116	**Robelle Consulting Ltd.** (888) 762-3553 or (604) 582-1700	89
102	**St. Paul Software** (612) 603-4400/Fax: (612) 603-4403	75
180	**Software Licensing Corp.** (800) 831-0882 or (702) 832-0881	95
168	**Sound Ideas of America** (800) 633-5250 or (717) 688-9511	9
166	**Straightline** (800) 458-1273/Fax: (206) 865-8314	53
141	**Stratus** (800) 486-9194	3
78	**Technical & Scientific Application** (800) 422-4872/Fax: (713) 935-1500	71
49	**Ted Dasher & Associates** (800) 638-4833/Fax: (205) 591-1108	93
47	**The Esys Corporation** (403) 424-4922/Fax: (403) 424-4925	75
106	**Tracer Technologies, Inc.** (301) 977-9400/Fax: (301) 977-5222	49
40	**VESOFT, Inc.** (310) 282-0420/Fax: (310) 785-9566	99
121	**Workstation Direct** (888) 622-1600/Fax: (800) 371-2657	77
132	**Xess Corp.** (800) 961-7840 or (919) 387-0076	25

IBM





IBM software for HP-UX

Cut superhuman tasks down to size.

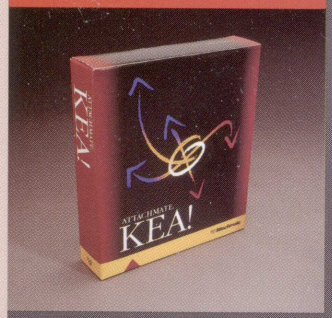
Michel Lotito of Grenoble, France was known as Monsieur Mangetout (" Mr. Eat-Everything"). Among the delectables he consumed were chandeliers, TV sets and an entire Cessna light aircraft. Now there's an appetite for big challenges. Which also describes IBM's DB2 for HP-UX. A great majority of the Fortune 500 already use DB2 to manage vast, complex databases. Today, this scalable solution is ready to go to work in your HP-UX environment – to help keep data safe, secure and easily accessible. In combination with IBM's Net.Data for HP-UX, DB2 can also help you deploy databases in dynamic Web applications. These are just two powerful and proven IBM solutions now available for HP-UX. Others include MQSeries, ADSM, Lotus Domino, Transaction Server and more. For further information – including opportunities for resellers – contact your HP distributor or visit the IBM Software for HP-UX Website at **www.software.ibm.com/hp**

CIRCLE 160 ON READER SERVICE CARD

IBM, MQSeries, DB2, Net.Data, ADSM (ADSTAR Distributed Storage Manager), and Transaction Server are registered trademarks of International Business Machines Corporation. Lotus Domino is a registered trademark of Lotus Corporation. All other company and/or product names are trademarks of their respective companies. ©1997 IBM Corp. All rights reserved. Reprinted from the Guinness Book of Records, 1997©. Permission granted by Guinness Media Inc.

KEA! 420 & 340

FOR VAX/UNIX



Choose the wrong connectivity software and it could reflect poorly on you.



FEATURES:

- *KEA! 420 for DEC® VT420 emulation*
- *KEA! 340 for DEC® VT340 graphics and VT420 emulation*
- *Supports DDE, OLE, HLLAPI, WinHLLAPI and Attachmate's Enterprise Access Library for automation and integration*
- *Fully customizable menus and SmartPads*
- *Macro language with macro recorder*
- *Multiple sessions*

New in Version 5.0:

- *One CD for Windows®, Windows 95 and Windows NT®*
- *Configuration wizard*
- *TrueType font*
- *Configurable toolbar*
- *Drivers for LK450 and Powerstation™ keyboards in Windows 95*

Call 1-800-933-6751

www.attachmate.com/ad/hpv.htm

Source Code: 630.02

Trade-up from your current VT connectivity software.

$100

Don't risk ruining your image. Choose KEA!,™ the complete solution for integrating VAX® and UNIX® host information with PC applications. Feature-rich and user-friendly, KEA! lets you create transparent links between Microsoft® Windows® documents and host documents.

CIRCLE 2 ON READER SERVICE CARD

Attachmate is a registered trademark and KEA! and Powerstation are trademarks of Attachmate Corporation. Microsoft Windows and Windows NT are registered trademarks of Microsoft Corporation. DEC and VAX are registered trademarks of Digital Equipment Corporation. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. TrueType is a trademark of Apple computer, Inc. All other trademarks or registered trademarks are the property of their respective owners.

Attachmate™